

microunity

Zeus System Architecture

COPYRIGHT 1998 MICROUNITY SYSTEMS ENGINEERING, INC. ALL RIGHTS RESERVED.



MicroUnity

Craig Hansen
Chief Architect

MicroUnity Systems Engineering, Inc.
475 Potrero Avenue
Sunnyvale, CA 94086.4118
Phone: 408.734.8100
Fax: 408.734.8136
email: craig@microunity.com
<http://www.microunity.com>

Contents

Contents	2
Tables and Figures	5
Introduction	6
Conformance	7
Mandatory and Optional Areas	7
Upward-compatible Modifications	7
Promotion of Optional Features	7
Unrestricted Physical Implementation	8
Draft Version	8
Common Elements	9
Notation	9
Bit ordering	10
Memory	10
Byte	10
Byte ordering	10
Memory read/load semantics	11
Memory write/store semantics	11
Data	11
Fixed-point Data	12
Address	14
Floating-point Data	14
Zeus Processor	17
Architectural Framework	17
Interfaces and Block Diagram	17
Instruction	17
Assembler Syntax	17
Instruction Structure	18
Gateway	18
User State	19
General Registers	19
Program Counter	20
Privilege Level	20
Program Counter and Privilege Level	20
System state	20
Fixed-point	21
Load and Store	21
Branch	21
Addressing Operations	22
Execution Operations	22
Floating-point	22
Branch Conditionally	22
Compare-set	23
Arithmetic Operations	23
Rounding and exceptions	24
NaN handling	25
Floating-point functions	26
Digital Signal Processing	35
Data-handling Operations	36
Arithmetic Operations	39
Galois Field Operations	39
Software Conventions	40
Register Usage	40
Procedure Calling Conventions	40
System and Privileged Library Calls	44
Instruction Scheduling	47
Separate Addressing from Execution	47
Software Pipeline	47
Multiple Issue	47
Functional Unit parallelism	47
Latency	48
Pipeline Organization	49
Classical Pipeline Structures	49
Superstring Pipeline	50
Superspring Pipeline	51
Superthread Pipeline	52
Simultaneous Multithreading	53
Branch/fetch Prediction	54
Additional Load and Execute Resources	55
Result Forwarding	55
Instruction Set	57
Major Operation Codes	58
Minor Operation Codes	59
General Forms	63
Instruction Fetch	64
Perform Exception	65
Instruction Decode	65
Always Reserved	72
Address	73
Address Compare	76
Address Copy Immediate	79
Address Immediate	80
Address Immediate Reversed	83
Address Reversed	86
Address Shift Left Immediate Add	89
Address Shift Left Immediate Subtract	90
Address Shift Immediate	91
Address Ternary	93
Branch	94
Branch Back	95
Branch Barrier	97
Branch Conditional	98
Branch Conditional Floating-Point	101
Branch Conditional Visibility Floating-Point	103
Branch Down	105
Branch Gateway	106
Branch Halt	110
Branch Hint	111
Branch Hint Immediate	112
Branch Immediate	113
Branch Immediate Link	114
Branch Link	115
Load	117
Load Immediate	120
Store	123
Store Double Compare Swap	126
Store Immediate	128
Store Immediate Inplace	131
Store Inplace	133

Group Add	135	Level One Cache	332
Group Add Halve	138	Level One Cache Stress Control	342
Group Boolean	141	Level One Cache Redundancy	342
Group Compare	148	Memory Attributes	343
Group Compare Floating-point	154	Cache Control	344
Group Copy Immediate	157	Cache Coherence	347
Group Immediate	159	Strong Ordering	348
Group Immediate Reversed	163	Victim Selection	349
Group Inplace	168	Detail Access	352
Group Reversed	170	Micro Translation Buffer	354
Group Reversed Floating-point	176	Block Translation Buffer	357
Group Shift Left Immediate Add	179	Program Translation Buffer	357
Group Shift Left Immediate Subtract	181	Global Virtual Cache	358
Group Subtract Halve	183	Memory Interface	358
Group Ternary	186	Microarchitecture	358
Crossbar	187	Snoop	359
Crossbar Extract	192	Load	359
Crossbar Field	196	Store	360
Crossbar Field Inplace	201	Memory	361
Crossbar Inplace	204	Bus interface	363
Crossbar Short Immediate	206	Motherboard Chipsets	363
Crossbar Short Immediate Inplace	211	Pinout	364
Crossbar Shuffle	213	Pin summary	364
Crossbar Swizzle	219	Electrical Specifications	368
Crossbar Ternary	220	Bus Control Register	373
Ensemble	221	Emulator signals	374
Ensemble Convolve Extract Immediate	225	A20M#	375
Ensemble Convolve Floating-point	232	INIT	375
Ensemble Extract	236	INTR	375
Ensemble Extract Immediate	244	NMI	375
Ensemble Extract Immediate Inplace	251	SMI#	376
Ensemble Floating-point	258	STPCLK#	376
Ensemble Inplace	261	IGNNE#	376
Ensemble Inplace Floating-point	264	Emulator output signals	377
Ensemble Reversed Floating-point	267	Bus snooping	377
Ensemble Ternary	269	Locked cycles	377
Ensemble Ternary Floating-point	272	Locked synchronization instruction	377
Ensemble Unary	274	Locked sequences of bus transactions	378
Ensemble Unary Floating-point	276	Sampled at Reset	378
Wide Multiply Matrix	283	Sampled per Clock	378
Wide Multiply Matrix Extract	288	Bus Access	379
Wide Multiply Matrix Extract Immediate	295	Other bus cycles	381
Wide Multiply Matrix Floating-point	303	Special cycles	381
Wide Multiply Matrix Galois	308	I/O cycles	382
Wide Switch	311	Events and Threads	384
Wide Translate	313	Ephemeral Program State	387
Memory Management	316	Event Register	388
Overview	316	Event Mask	390
Local Translation Buffer	317	Exceptions:	392
Global Translation Buffer	321	GlobalTBMiss Handler	392
GTB Registers	324	Exceptions in detail	395
Address Generation	326	Reserved Instruction	395
Memory Banks	327	Access Disallowed by virtual address	396
Program Microcache	328	Access disallowed by tag	396
Wide Microcache	329	Access detail required by tag	396
Level Zero Cache	330	Access disallowed by global TB	396
Structure	331	Access detail required by global TB	396

Global TB miss	397	Reset state	402
Access disallowed by local TB	397	Start Address.....	402
Access detail required by local TB.....	397	Internal ROM Code.....	403
Local TB miss.....	398	Memory and Devices	404
Floating-point arithmetic	398	Physical Memory Map.....	404
Fixed-point arithmetic.....	398	Architecture Description Register	407
Reset and Error Recovery.....	399	Status Register	407
Reset.....	399	Control Register	409
Power-on Reset.....	399	Clock.....	412
Bus Reset.....	399	Clock Cycle.....	412
Control Register Reset.....	400	Clock Event.....	412
Meltdown Detected Reset.....	400	Clock Watchdog	412
Double Check Reset.....	400	Tally.....	414
Machine Check	400	Tally Counter.....	414
Parity or Uncorrectable Error in Cache..	401	Tally Control	415
Parity or Communications Error in Bus.	401	Thread Register	417
Watchdog Timeout Error.....	402	Index	418
Event Thread Exception.....	402		

Tables and Figures

descriptive notation	9
compare-branch relations	23
compare-set relations	23
32-bit 2-way deal	36
16-bit 4-way deal	37
16-bit 2-way shuffle	37
16-bit 4-way shuffle	37
16-bit reverse	38
Compress 32 bits to 16, with 4-bit right shift ...	38
Expand 16 bits to 32, with 4-bit left shift.....	39
register usage.....	40
Alignment within dp region.....	42
Gateway with pointers to code and data spaces	45
canonical pipeline.....	49
canonical pipeline.....	49
superscalar pipeline.....	50
superpipelined pipeline.....	50
Superstring pipeline	51
Superspring pipeline	52
Superspring pipeline	52
Superthread pipeline	53
Superthread pipeline	53
major operation code field values	58
minor operation code field values for A.MINOR	59
minor operation code field values for B.MINOR	59
minor operation code field values for L.MINOR	59
minor operation code field values for S.MINOR	60
minor operation code field values for G.size....	60
minor operation code field values for XSHIFTI	60
minor operation code field values for XSHIFT	60
minor operation code field values for E.size....	60
minor operation code field values for W.MINOR.L or W.MINOR.B	61
minor operation code field values for EMULXI, EMULXIU, EMULXIM, EMULXIC, EMULADDXI, EMULADDXIU,	
EMULADDXIM, EMULADDXIC, ECONXIL, ECONXIB, ECONXIUL, ECONXIUB, ECONXIML, ECONXIMB, ECONXICL, ECONXICB, EEXTRACTI, EEXTRACTIU, WMULMATXIUL, WMULMATXIUB, WMULMATXIML, WMULMATXIMB, WMULMATXICL, and WMULMATXICB,...	61
operand size field values for G.COPY.I, G.AND.I, G.NAND.I, G.NOR.I, G.OR.I, G.XOR.I, G.ADD.I, G.ADD.I.O, G.ADD.I.UO, G.SET.AND.E.I, G.SET.AND.NE.I, G.SET.E.I, G.SET.GE.I, G.SET.L.I, G.SET.NE.I, G.SET.GE.I.U, G..SET.L.I.U, G.SUB.I,	
G.SUB.I.O, G.SUB.I.UO	62
unary operation code field values for E.UNARY.size	62

compare operation code field values for A.COM.op and G.COM.op.size.....	63
Branch gateway	107
Crossbar extract	193
Crossbar merge extract	193
4-way shuffle bytes within hexlet	214
4-way shuffle bytes within triclet.....	214
Ensemble convolve extract immediate doublets	229
Ensemble convolve extract immediate complex doublets.....	230
Ensemble convolve floating-point half little-endian	233
Ensemble convolve complex floating-point half little-endian	234
Ensemble complex multiply extract doublets..	239
Ensemble scale add extract doublets	239
Ensemble complex scale add extract doublets	240
Ensemble extract	240
Ensemble merge extract	241
Ensemble multiply extract immediate doublets	247
Ensemble multiply extract immediate doublets	247
Ensemble multiply extract immediate complex doublets.....	248
Ensemble multiply extract immediate complex doublets.....	248
Ensemble multiply add extract immediate doublets	253
Ensemble multiply add extract immediate doublets	254
Ensemble multiply add extract immediate complex doublets.....	255
Ensemble multiply add extract immediate complex doublets.....	256
Ensemble multiply Galois field bytes	270
Wide multiply matrix	285
Wide multiply matrix complex.....	285
Wide multiply extract matrix doublets.....	290
Wide multiply extract matrix complex doublets	291
Wide multiply matrix extract immediate doublets	299
Wide multiply matrix extract immediate complex doublets.....	300
Wide multiply matrix floating-point half.....	304
Wide multiply matrix complex floating-point half	305
Wide multiply matrix Galois	309
memory management organization	316
local virtual address space specifiers	318
machine check errors	401
tally control field interpretation	416

Introduction

MicroUnity's Zeus Architecture describes general-purpose processor, memory, and interface subsystems, organized to operate at the enormously high bandwidth rates required for broadband applications.

The Zeus processor performs integer, floating point, signal processing and non-linear operations such as Galois field, table lookup and bit switching on data sizes from 1 bit to 128 bits. Group or SIMD (single instruction multiple data) operations sustain external operand bandwidth rates up to 512 bits (i.e., up to four 128-bit operand groups) per instruction even on data items of small size. The processor performs ensemble operations such as convolution that maintain full intermediate precision with aggregate internal operand bandwidth rates up to 20,000 bits per instruction. The processor performs wide operations such as crossbar switch, matrix multiply and table lookup that use caches embedded in the execution units themselves to extend operands to as much as 32768 bits. All instructions produce at most a single 128-bit register result, source at most three 128-bit registers and are free of side effects such as the setting of condition codes and flags. The instruction set design carries the concept of streamlining beyond Reduced Instruction Set Computer (RISC) architectures, to simplify implementations that issue several instructions per machine cycle.

The Zeus memory subsystem provides 64-bit virtual and physical addressing for UNIX, Mach, and other advanced OS environments. Separate address instructions enable the division of the processor into decoupled access and execution units, to reduce the effective latency of memory to the pipeline. The Zeus cache supplies the high data and instruction issue rates of the processor, and supports coherency primitives for scalable multiprocessors. The memory subsystem includes mechanisms for sustaining high data rates not only in block transfer modes, but also in non-unit stride and scattered access patterns.

The Zeus interface subsystem is designed to match industry-standard "Socket 7" protocols and pin-outs. In this way, Zeus can make use of the immense infrastructure of the PC for building low-cost systems. The interface subsystem is modular, and can be replaced with appropriate protocols and pin-outs for lower-cost and higher-performance systems.

The goal of the Zeus architecture is to integrate these processor, memory, and interface capabilities with optimal simplicity and generality. From the software perspective, the entire machine state consists of a program counter, a single bank of 64 general-purpose 128-bit registers, and a linear byte-addressed shared memory space with mapped interface registers. All interrupts and exceptions are precise, and occur with low overhead.

This document is intended for Zeus software and hardware developers alike, and defines the interface at which their designs must meet. Zeus pursues the most efficient tradeoffs between hardware and software complexity by making all processor, memory, and interface resources directly accessible to high-level language programs.

Conformance

To ensure that Zeus systems may freely interchange data, user-level programs, system-level programs and interface devices, the Zeus system architecture reaches above the processor level architecture.

Mandatory and Optional Areas

A computer system conforms to the requirements of the Zeus System Architecture if and only if it implements all the specifications described in this document and other specifications included by reference. Conformance to the specification is mandatory in all areas, including the instruction set, memory management system, interface devices and external interfaces, and bootstrap ROM functional requirements, except where explicit options are stated.

Optional areas include:

- Number of processor threads
- Size of first-level cache memories
- Existence of a second-level cache
- Size of second-level cache memory
- Size of system-level memory
- Existence of certain optional interface device interfaces

Upward-compatible Modifications

From time to time, MicroUnity may modify the architecture in an upward-compatible manner, such as by the addition of new instructions, definition of reserved bits in system state, or addition of new standard interfaces. Such modifications will be added as options, so that designs that conform to this version of the architecture will conform to future, modified versions.

Additional devices and interfaces, not covered by this standard may be added in specified regions of the physical memory space, provided that system reset places these devices and interfaces in an inactive state that does not interfere with the operation of software that runs in any conformant system. The software interface requirements of any such additional devices and interfaces must be made as widely available as this architecture specification.

Promotion of Optional Features

It is most strongly recommended that such optional instructions, state or interfaces be implemented in all conforming designs. Such implementations enhance the value of the features in particular and the architecture as a whole by broadening the set of implementations over which software may depend upon the presence of these features.

Implementations that fail to implement these features may encounter unacceptable levels of overhead when attempting to emulate the features by exception handlers or use of virtual

memory. This is a particular concern when involved in code that has real-time performance constraints.

In order that upward-compatible optional extensions of the original Zeus system architecture may be relied upon by system and application software, MicroUnity may upon occasion promote optional features to mandatory conformance for implementations designed or produced after a suitable delay upon such notification by publication of future version of the specification.

Unrestricted Physical Implementation

Nothing in this specification should be construed to limit the implementation choices of the conforming system beyond the specific requirements stated herein. In particular, a computer system may conform to the Zeus System Architecture while employing any number of components, dissipate any amount of heat, require any special environmental facilities, or be of any physical size.

Draft Version

This document is a draft version of the architectural specification. In this form, conformance to this document may not be claimed or implied. MicroUnity may change this specification at any time, in any manner, until it has been declared final. When this document has been declared final, the only changes will be to correct bugs, defects or deficiencies, and to add upward-compatible optional extensions.

Common Elements

Notation

The descriptive notation used in this document is summarized in the table below:

$x + y$	two's complement addition of x and y . Result is the same size as the operands, and operands must be of equal size.
$x - y$	two's complement subtraction of y from x . Result is the same size as the operands, and operands must be of equal size.
$x * y$	two's complement multiplication of x and y . Result is the same size as the operands, and operands must be of equal size.
x / y	two's complement division of x by y . Result is the same size as the operands, and operands must be of equal size.
$x \& y$	bitwise and of x and y . Result is same size as the operands, and operands must be of equal size.
$x \mid y$	bitwise or of x and y . Result is same size as the operands, and operands must be of equal size.
$x ^ y$	bitwise exclusive-or of x and y . Result is same size as the operands, and operands must be of equal size.
$\sim x$	bitwise inversion of x . Result is same size as the operand.
$x = y$	two's complement equality comparison between x and y . Result is a single bit, and operands must be of equal size.
$x \neq y$	two's complement inequality comparison between x and y . Result is a single bit, and operands must be of equal size.
$x < y$	two's complement less than comparison between x and y . Result is a single bit, and operands must be of equal size.
$x \geq y$	two's complement greater than or equal comparison between x and y . Result is a single bit, and operands must be of equal size.
\sqrt{x}	floating-point square root of x
$x \amalg y$	concatenation of bit field x to left of bit field y
x^y	binary digit x repeated, concatenated y times. Size of result is y .
x_y	extraction of bit y (using little-endian bit numbering) from value x . Result is a single bit.
$x_{y..z}$	extraction of bit field formed from bits y through z of value x . Size of result is $y-z+1$; if $z>y$, result is an empty string.
$x?y:z$	value of y , if x is true, otherwise value of z . Value of x is a single bit.
$x \leftarrow y$	bitwise assignment of x to value of y
S_n	signed, two's complement, binary data format of n bytes
U_n	unsigned binary data format of n bytes
F_n	floating-point data format of n bytes


descriptive notation

Bit ordering

The ordering of bits in this document is always little-endian, regardless of the ordering of bytes within larger data structures. Thus, the least-significant bit of a data structure is always labeled 0 (zero), and the most-significant bit is labeled as the data structure size (in bits) minus one.


Memory

Zeus memory is an array of 2^{64} bytes, without a specified byte ordering, which is physically distributed among various components.



Byte


A byte is a single element of the memory array, consisting of 8 bits:




Byte ordering

Larger data structures are constructed from the concatenation of bytes in either little-endian or big-endian byte ordering. A memory access of a data structure of size s at address i is formed from memory bytes at addresses i through i+s-1. Unless otherwise specified, there is no specific requirement of alignment: it is not generally required that i be a multiple of s. Aligned accesses are preferred whenever possible, however, as they will often require one fewer processor or memory clock cycle than unaligned accesses.

With little-endian byte ordering, the bytes are arranged as:



With big-endian byte ordering, the bytes are arranged as:



Zeus memory is byte-addressed, using either little-endian or big-endian byte ordering. For consistency with the bit ordering, and for compatibility with x86 processors, Zeus uses little-endian byte ordering when an ordering must be selected. Zeus load and store instructions are available for both little-endian and big-endian byte ordering. The selection of byte ordering is dynamic, so that little-endian and big-endian processes, and even data structures within a process, can be intermixed on the processor.

Memory read/load semantics

Zeus memory, including memory-mapped registers, must conform to the following requirements regarding side-effects of read or load operations:

A memory read must have no side-effects on the contents of the addressed memory nor on the contents of any other memory.

Memory write/store semantics

Zeus memory, including memory-mapped registers, must conform to the following requirements regarding side-effects of read or load operations:

A memory write must affect the contents of the addressed memory so that a memory read of the addressed memory returns the value written, and so that a memory read of a portion of the addressed memory returns the appropriate portion of the value written.

A memory write may affect or cause side-effects on the contents of memory not addressed by the write operation, however, a second memory write of the same value to the same address must have no side-effects on any memory; memory write operations must be idempotent.

Zeus store instructions that are weakly ordered may have side-effects on the contents of memory not addressed by the store itself; subsequent load instructions which are also weakly ordered may or may not return values which reflect the side-effects.

Data

Zeus provides eight-byte (64-bit) virtual and physical address sizes, and eight-byte (64-bit) and sixteen-byte (128-bit) data path sizes, and uses fixed-length four-byte (32-bit) instructions. Arithmetic is performed on two's-complement or unsigned binary and ANSI/IEEE standard 754-1985 conforming binary floating-point number representations.

Fixed-point DataBit


A bit is a primitive data element:

Peck


A peck is the catenation of two bits:

Nibble


A nibble is the catenation of four bits:

Byte


A byte is the catenation of eight bits, and is a single element of the memory array:

Doublet

A doublet is the catenation of 16 bits, and is the catenation of two bytes:


Quadlet

A quadlet is the catenation of 32 bits, and is the catenation of four bytes:




Octlet

An octlet is the catenation of 64 bits, and is the catenation of eight bytes:


Hexlet

A hexlet is the catenation of 128 bits, and is the catenation of sixteen bytes:



Triclet

A triclet is the catenation of 256 bits, and is the catenation of thirty-two bytes:

Address

Zeus addresses, both virtual addresses and physical addresses, are octlet quantities.

Floating-point Data

Zeus's floating-point formats are designed to satisfy ANSI/IEEE standard 754-1985: Binary Floating-point Arithmetic. Standard 754 leaves certain aspects to the discretion of implementers: additional precision formats, encoding of quiet and signaling NaN values, details of production and propagation of quiet NaN values. These aspects are detailed below.

Zeus adds additional half-precision and quad-precision formats to standard 754's single-precision and double-precision formats. Zeus's double-precision satisfies standard 754's

precision requirements for a single-extended format, and Zeus's quad-precision satisfies standard 754's precision requirements for a double-extended format.

Each precision format employs fields labeled **s** (sign), **e** (exponent), and **f** (fraction) to encode values that are (1) NaN: quiet and signaling, (2) infinities: $(-1)^{\text{s}\infty}$, (3) normalized numbers: $(-1)^{\text{s}2^{\text{e}-\text{bias}}}(1.\text{f})$, (4) denormalized numbers: $(-1)^{\text{s}2^{1-\text{bias}}}(0.\text{f})$, and (5) zero: $(-1)^{\text{s}0}$.

Quiet NaN values are denoted by any sign bit value, an exponent field of all one bits, and a non-zero fraction with the most significant bit set. Quiet NaN values generated by default exception handling of standard operations have a zero sign bit, an exponent field of all one bits, a fraction field with the most significant bit set, and all other bits cleared.

Signaling NaN values are denoted by any sign bit value, an exponent field of all one bits, and a non-zero fraction with the most significant bit cleared.

Infinite values are denoted by any sign bit value, an exponent field of all one bits, and a zero fraction field.

Normalized number values are denoted by any sign bit value, an exponent field that is not all one bits or all zero bits, and any fraction field value. The numeric value encoded is $(-1)^{\text{s}2^{\text{e}-\text{bias}}}(1.\text{f})$. The bias is equal the value resulting from setting all but the most significant bit of the exponent field, half: 15, single: 127, double: 1023, and quad: 16383.

Denormalized number values are denoted by any sign bit value, an exponent field that is all zero bits, and a non-zero fraction field value. The numeric value encoded is $(-1)^{\text{s}2^{1-\text{bias}}}(0.\text{f})$.

Zero values are denoted by any sign bit value, and exponent field that is all zero bits, and a fraction field that is all zero bits. The numeric value encoded is $(-1)^{\text{s}0}$. The distinction between +0 and -0 is significant in some operations.


Half-precision Floating-point

Zeus half precision uses a format similar to standard 754's requirements, reduced to a 16-bit overall format. The format contains sufficient precision and exponent range to hold a 12-bit signed integer.


1514			109	0
s	e	f		
1	5	10		

Single-precision Floating-point


Zeus single precision satisfies standard 754's requirements for "single."

Double-precision Floating-point

Zeus double precision satisfies standard 754's requirements for "double."

Quad-precision Floating-point

Zeus quad precision satisfies standard 754's requirements for "double extended," but has additional fraction precision to use 128 bits.



Zeus Processor

MicroUnity's Zeus processor provides the general-purpose, high-bandwidth computation capability of the Zeus system. Zeus includes high-bandwidth data paths, register files, and a memory hierarchy. Zeus's memory hierarchy includes on-chip instruction and data memories, instruction and data caches, a virtual memory facility, and interfaces to external devices. Zeus's interfaces in the initial implementation are solely the "Super Socket 7" bus, but other implementations may have different or additional interfaces.

Architectural Framework

The Zeus architecture defines a compatible framework for a family of implementations with a range of capabilities. The following implementation-defined parameters are used in the rest of the document in boldface. The value indicated is for MicroUnity's first Zeus implementation.

Parameter	Interpretation	Value	Range of legal values
T	number of execution threads	4	$1 \leq T \leq 31$
CE	\log_2 cache blocks in first-level cache	9	$0 \leq CE \leq 31$
CS	\log_2 cache blocks in first-level cache set	2	$0 \leq CS \leq 4$
CT	existence of dedicated tags in first-level cache	1	$0 \leq CT \leq 1$
LE	\log_2 entries in local TB	0	$0 \leq LE \leq 3$
LB	Local TB based on base register	1	$0 \leq LB \leq 1$
GE	\log_2 entries in global TB	7	$0 \leq GE \leq 15$
GT	\log_2 threads which share a global TB	1	$0 \leq GT \leq 3$

Interfaces and Block Diagram

The first implementation of Zeus uses "socket 7" protocols and pinouts.

Instruction

Assembler Syntax

Instructions are specified to Zeus assemblers and other code tools (assemblers) in the syntax of an instruction mnemonic (operation code), then optionally white space (blanks or tabs) followed by a list of operands.

The instruction mnemonics listed in this specification are in upper case (capital) letters, assemblers accept either upper case or lower case letters in the instruction mnemonics. In


this specification, instruction mnemonics contain periods (“.”) to separate elements to make them easier to understand; assemblers ignore periods within instruction mnemonics. The instruction mnemonics are designed to be parsed uniquely without the separating periods.

If the instruction produces a register result, this operand is listed first. Following this operand, if there are one or more source operands, is a separator which may be a comma (“,”), equal (“=”), or at-sign (“@”). The equal separates the result operand from the source operands, and may optionally be expressed as a comma in assembler code. The at-sign indicates that the result operand is also a source operand, and may optionally be expressed as a comma in assembler code. If the instruction specification has an equal-sign, an at-sign in assembler code indicates that the result operand should be repeated as the first source operand (for example, “A.ADD.I r4@5” is equivalent to “A.ADD.I r4=r4,5”). Commas always separate the remaining source operands.

The result and source operands are case-sensitive; upper case and lower case letters are distinct. Register operands are specified by the names r0 (or r00) through r63 (a lower case “r” immediately followed by a one or two digit number from 0 to 63), or by the special designations of “lp” for “r0,” “dp” for “r1,” “fp” for “r62,” and “sp” for “r63.” Integer-valued operands are specified by an optional sign (-) or (+) followed by a number, and assemblers generally accept a variety of integer-valued expressions.


Instruction Structure

A Zeus instruction is specifically defined as a four-byte structure with the little-endian ordering shown below. It is different from the quadlet defined above because the placement of instructions into memory must be independent of the byte ordering used for data structures. Instructions must be aligned on four-byte boundaries; in the diagram below, i must be a multiple of 4.




Gateway

A Zeus gateway is specifically defined as an 8-byte structure with the little-endian ordering shown below. A gateway contains a code address used to securely invoke a system call or procedure at a higher privilege level. Gateways are marked by protection information specified in the TB. Gateways must be aligned on 8-byte boundaries; in the diagram below, i must be a multiple of 8.



The gateway contains two data items within its structure, a code address and a new privilege level:




The virtual memory system can be used to designate a region of memory as containing gateways. Other data may be placed within the gateway region, provided that if an attempt is made to use the additional data as a gateway, that security cannot be violated. For example, 64-bit data or stack pointers which are aligned to at least 4 bytes and are in little-endian byte order have pl=0, so that the privilege level cannot be raised by attempting to use the additional data as a gateway.

User State

The user state consists of hardware data structures that are accessible to all conventional compiled code. The Zeus user state is designed to be as regular as possible, and consists only of the general registers, the program counter, and virtual memory. There are no specialized registers for condition codes, operating modes, rounding modes, integer multiply/divide, or floating-point values.

General Registers

Zeus user state includes 64 general registers. All are identical; there is no dedicated zero-valued register, and there are no dedicated floating-point registers.



Some Zeus instructions have 64-bit register operands. These operands are sign-extended to 128 bits when written to the register file, and the low-order 64 bits are chosen when read from the register file.

Definition

```
def val ← RegRead[rn, size]
  case size of
    64:
      val ← REG[rn]63..0
    128:
```

```


        val ← REG[rn]
      endcase
enddef

def RegWrite(rn, size, val)
  case size of
    64:
      REG[rn] ← val6364 || val63..0
    128:
      REG[rn] ← val127..0
  endcase
enddef

```

Program Counter

The program counter contains the address of the currently executing instruction. This register is implicitly manipulated by branch instructions, and read by branch instructions that save a return address in a general register.



Privilege Level

The privilege level register contains the privilege level of the currently executing instruction. This register is implicitly manipulated by branch gateway and branch down instructions, and read by branch gateway instructions that save a return address in a general register.



Program Counter and Privilege Level

The program counter and privilege level may be packed into a single octlet. This combined data structure is saved by the Branch Gateway instruction and restored by the Branch Down instruction.



System state

The system state consists of the facilities not normally used by conventional compiled code. These facilities provide mechanisms to execute such code in a fully virtual environment. All system state is memory mapped, so that it can be manipulated by compiled code.

Fixed-point

Zeus provides load and store instructions to move data between memory and the registers, branch instructions to compare the contents of registers and to transfer control from one code address to another, and arithmetic operations to perform computation on the contents of registers, returning the result to registers.

Load and Store

The load and store instructions move data between memory and the registers. When loading data from memory into a register, values are zero-extended or sign-extended to fill the register. When storing data from a register into memory, values are truncated on the left to fit the specified memory region.

Load and store instructions that specify a memory region of more than one byte may use either little-endian or big-endian byte ordering: the size and ordering are explicitly specified in the instruction. Regions larger than one byte may be either aligned to addresses that are an even multiple of the size of the region or of unspecified alignment: alignment checking is also explicitly specified in the instruction.

Load and store instructions specify memory addresses as the sum of a base general register and the product of the size of the memory region and either an immediate value or another general register. Scaling maximizes the memory space which can be reached by immediate offsets from a single base general register, and assists in generating memory addresses within iterative loops. Alignment of the address can be reduced to checking the alignment of the first general register.

The load and store instructions are used for fixed-point data as well as floating-point and digital signal processing data; Zeus has a single bank of registers for all data types.

Swap instructions provide multithread and multiprocessor synchronization, using indivisible operations: add-swap, compare-swap, multiplex-swap, and double-compare-swap. A store-multiplex operation provides the ability to indivisibly write to a portion of an octlet. These instructions always operate on aligned octlet data, using either little-endian or big-endian byte ordering.

Branch

The fixed-point compare-and-branch instructions provide all arithmetic tests for equality and inequality of signed and unsigned fixed-point values. Tests are performed either between two operands contained in general registers, or on the bitwise and of two operands. Depending on the result of the compare, either a branch is taken, or not taken. A taken branch causes an immediate transfer of the program counter to the target of the branch, specified by a 12-bit signed offset from the location of the branch instruction. A non-taken branch causes no transfer; execution continues with the following instruction.

Other branch instructions provide for unconditional transfer of control to addresses too distant to be reached by a 12-bit offset, and to transfer to a target while placing the location

following the branch into a register. The branch through gateway instruction provides a secure means to access code at a higher privilege level, in a form similar to a normal procedure call.

Addressing Operations

A subset of general fixed-point arithmetic operations is available as addressing operations. These include add, subtract, Boolean, and simple shift operations. These addressing operations may be performed at a point in the Zeus processor pipeline so that they may be completed prior to or in conjunction with the execution of load and store operations in a "superspring" pipeline in which other arithmetic operations are deferred until the completion of load and store operations.

Execution Operations

Many of the operations used for Digital Signal Processing (DSP), which are described in greater detail below, are also used for performing simple scalar operations. These operations perform arithmetic operations on values of 8-, 16-, 32-, 64-, or 128-bit sizes, which are right-aligned in registers. These execution operations include the add, subtract, boolean and simple shift operations which are also available as addressing operations, but further extend the available set to include three-operand add/subtract, three-operand boolean, dynamic shifts, and bit-field operations.

Floating-point

Zeus provides all the facilities mandated and recommended by ANSI/IEEE standard 754-1985: Binary Floating-point Arithmetic, with the use of supporting software.

Branch Conditionally

The floating-point compare-and-branch instructions provide all the comparison types required and suggested by the IEEE floating-point standard. These floating-point comparisons augment the usual types of numeric value comparisons with special handling for NaN (not-a-number) values. A NaN value compares as "unordered" with respect to any other value, even that of an identical NaN value.

Zeus floating-point compare-branch instructions do not generate an exception on comparisons involving quiet or signaling NaN values. If such exceptions are desired, they can be obtained by combining the use of a floating-point compare-set instruction, with either a floating-point compare-branch instruction on the floating-point operands or a fixed-point compare-branch on the set result.

Because the less and greater relations are anti-commutative, one of each relation that differs from another only by the replacement of an L with a G in the code can be removed by reversing the order of the operands and using the other code. Thus, an L relation can be used in place of a G relation by swapping the operands to the compare-branch or compare-set instruction.

No instructions are provided that branch when the values are unordered. To accomplish such an operation, use the reverse condition to branch over an immediately following unconditional branch, or in the case of an if-then-else clause, reverse the clauses and use the reverse condition.

The E relation can be used to determine the unordered condition of a single operand by comparing the operand with itself.

The following floating-point compare-branch relations are provided as instructions:

Mnemonic		Branch taken if values compare as:				Exception if	
code	C-like	Unord- ered	Greater	Less	Equal	unord- ered	invalid
E	==	F	F	F	T	no	no
LG	<>	F	T	T	F	no	no
L	<	F	F	T	F	no	no
GE	>=	F	T	F	T	no	no

compare-branch relations

Compare-set

The compare-set floating-point instructions provide all the comparison types supported as branch instructions. Zeus compare-set floating-point instructions may optionally generate an exception on comparisons involving quiet or signaling NaNs.

The following floating-point compare-set relations are provided as instructions:

Mnemonic		Result if values compare as:				Exception if	
code	C-like	Unord- ered	Greater	Less	Equal	unord- ered	invalid
E	==	F	F	F	T	no	no
LG	<>	F	T	T	F	no	no
L	<	F	F	T	F	no	no
GE	>=	F	T	F	T	no	no
EX	==	F	F	F	T	no	yes
LGX	<>	F	T	T	F	no	yes
LX	<	F	F	T	F	yes	yes
GEX	<=	F	T	F	T	yes	yes

compare-set relations

Arithmetic Operations

The basic operations supported in hardware are floating-point add, subtract, multiply, divide, square root and conversions among floating-point formats and between floating-point and binary integer formats.

Software libraries provide other operations required by the ANSI/IEEE floating-point standard.

The operations explicitly specify the precision of the operation, and round the result (or check that the result is exact) to the specified precision at the conclusion of each operation. Each of the basic operations splits operand registers into symbols of the specified precision and performs the same operation on corresponding symbols.

In addition to the basic operations, Zeus performs a variety of operations in which one or more products are summed to each other and/or to an additional operand. The instructions include a fused multiply-add (E.MUL.ADD.F), convolve (E.CON.F), matrix multiply (E.MUL.MAT.F), and scale-add (E.SCAL.ADD.F).

The results of these operations are computed as if the multiplies are performed to infinite precision, added as if in infinite precision, then rounded only once. Consequently, these operations perform these operations with no rounding of intermediate results that would have limited the accuracy of the result.

Rounding and exceptions

Rounding is specified within the instructions explicitly, to avoid explicit state registers for a rounding mode. Similarly, the instructions explicitly specify how standard exceptions (invalid operation, division by zero, overflow, underflow and inexact) are to be handled.¹

When no rounding is explicitly named by the instruction (default), round to nearest rounding is performed, and all floating-point exception signals cause the standard-specified default result, rather than a trap. When rounding is explicitly named by the instruction (N: nearest, Z: zero, F: floor, C: ceiling), the specified rounding is performed, and floating-point exception signals other than inexact cause a floating-point exception trap. When X (exact, or exception) is specified, all floating-point exception signals cause a floating-point exception trap, including inexact.

This technique assists the Zeus processor in executing floating-point operations with greater parallelism. When default rounding and exception handling control is specified in floating-point instructions, Zeus may safely retire instructions following them, as they are guaranteed not to cause data-dependent exceptions. Similarly, floating-point instructions with N, Z, F, or C control can be guaranteed not to cause data-dependent exceptions once the operands have been examined to rule out invalid operations, division by zero, overflow or underflow exceptions. Only floating-point instructions with X control, or when exceptions cannot be ruled out with N, Z, F, or C control need to avoid retiring following instructions until the final result is generated.

ANSI/IEEE standard 754-1985 specifies information to be given to trap handlers for the five floating-point exceptions. The Zeus architecture produces a precise exception, (The program counter points to the instruction that caused the exception and all register state is present) from which all the required information can be produced in software, as all source operand values and the specified operation are available.

¹U.S. Patent 5,812,439 describes this "Technique of incorporating floating point information into processor instructions."

ANSI/IEEE standard 754-1985 specifies a set of five “sticky-exception” bits, for recording the occurrence of exceptions that are handled by default. The Zeus architecture produces a precise exception for instructions with N, Z, F, or C control for invalid operation, division by zero, overflow or underflow exceptions and with X control for all floating-point exceptions, from which corresponding sticky-exception bits can be set. Execution of the same instruction with default control will compute the default result with round-to-nearest rounding. Most compound operations not specified by the standard are not available with rounding and exception controls.

NaN handling

ANSI/IEEE standard 754-1985 specifies that operations involving a signaling NaN or invalid operation shall, if no trap occurs and if a floating-point result is to be delivered, deliver a quiet NaN as its result. However, it fails to specify what quiet NaN value to deliver.

Zeus operations that produce a floating-point result and do not trap on invalid operations propagate signaling NaN values from operands to results, changing the signaling NaN values to quiet NaN values by setting the most significant fraction bit and leaving the remaining bits unchanged. Other causes of invalid operations produce the default quiet NaN value, where the sign bit is zero, the exponent field is all one bits, the most significant fraction bit is set and the remaining fraction bits are zero bits. For Zeus operations that produce multiple results catenated together, signaling NaN propagation or quiet NaN production is handled separately and independently for each result symbol.

ANSI/IEEE standard 754-1985 specifies that quiet NaN values should be propagated from operand to result by the basic operations. However, it fails to specify which of several quiet NaN values to propagate when more than one operand is a quiet NaN. In addition, the standard does not clearly specify how quiet NaN should be propagated for the multiple-operation instructions provided in Zeus. The standard does not specify the quiet NaN produced as a result of an operand being a signaling NaN when invalid operation exceptions are handled by default. The standard leaves unspecified how quiet and signaling NaN values are propagated through format conversions and the absolute-value, negate and copy operations. This section specifies these aspects left unspecified by the standard.

First of all, for Zeus operations that produce multiple results catenated together, quiet and signaling NaN propagation is handled separately and independently for each result symbol. A quiet or signaling NaN value in a single symbol of an operand causes only those result symbols that are dependent on that operand symbol's value to be propagated as that quiet NaN. Multiple quiet or signaling NaN values in symbols of an operand which influence separate symbols of the result are propagated independently of each other. Any signaling NaN that is propagated has the high-order fraction bit set to convert it to a quiet NaN.

For Zeus operations in which multiple symbols among operands upon which a result symbol is dependent are quiet or signaling NaNs, a priority rule will determine which NaN is propagated. Priority shall be given to the operand that is specified by a register definition at a lower-numbered (little-endian) bit position within the instruction (rb has priority over rc, which has priority over rd). In the case of operands which are catenated from two registers, priority shall be assigned based on the register which has highest priority (lower-numbered

bit position within the instruction). In the case of tie (as when the E.SCAL.ADD scaling operand has two corresponding NaN values, or when a E.MUL.CF operand has NaN values for both real and imaginary components of a value), the value which is located at a lower-numbered (little-endian) bit position within the operand is to receive priority. The identification of a NaN as quiet or signaling shall not confer any priority for selection – only the operand position, though a signaling NaN will cause an invalid operand exception.

The sign bit of NaN values propagated shall be complemented if the instruction subtracts or negates the corresponding operand or (but not and) multiplies it by or divides it by or divides it into an operand which has the sign bit set, even if that operand is another NaN. If a NaN is both subtracted and multiplied by a negative value, the sign bit shall be propagated unchanged.

For Zeus operations that convert between two floating-point formats (INFLATE and DEFLATE), NaN values are propagated by preserving the sign and the most-significant fraction bits, except that the most-significant bit of a signalling NaN is set and (for DEFLATE) the least-significant fraction bit preserved is combined, via a logical-or of all fraction bits not preserved. All additional fraction bits (for INFLATE) are set to zero.

For Zeus operations that convert from a floating-point format to a fixed-point format (SINK), NaN values produce zero values (maximum-likelihood estimate). Infinity values produce the largest representable positive or negative fixed-point value that fits in the destination field. When exception traps are enabled, NaN or Infinity values produce a floating-point exception. Underflows do not occur in the SINK operation, they produce -1, 0 or +1, depending on rounding controls.

For absolute-value, negate, or copy operations, NaN values are propagated with the sign bit cleared, complemented, or copied, respectively. Signalling NaN values cause the Invalid operation exception, propagating a quieted NaN in corresponding symbol locations (default) or an exception, as specified by the instruction.

Floating-point functions

The following functions are defined for use within the detailed instruction definitions in the following section. In these functions an internal format represents infinite-precision floating-point values as a four-element structure consisting of (1) s (sign bit): 0 for positive, 1 for negative, (2) t (type): NORM, ZERO, SNAN, QNAN, INFINITY, (3) e (exponent), and (4) f: (fraction). The mathematical interpretation of a normal value places the binary point at the units of the fraction, adjusted by the exponent: $(-1)^s \cdot (2^e) \cdot f$. The function F converts a packed IEEE floating-point value into internal format. The function PackF converts an internal format back into IEEE floating-point format, with rounding and exception control.

Definition

```
def eb ← ebits(prec) as
  case pref of
    16:
      eb ← 5
    32:
```

```

eb ← 8
64:
    eb ← 11
128:
    eb ← 15
endcase
enddef

def eb ← ebias(prec) as
    eb ← 0 11 1ebits(prec)-1
enddef

def fb ← fbits(prec) as
    fb ← prec - 1 - eb
enddef

def a ← F(prec, ai) as
    a.s ← aiprec-1
    ae ← aiprec-2..fbits(prec)
    af ← aifbits(prec)-1..0
    if ae = 1ebits(prec) then
        if af = 0 then
            a.t ← INFINITY
        elseif affbits(prec)-1 then
            a.t ← SNaN
            a.e ← -fbits(prec)
            a.f ← 1 11 affbits(prec)-2..0
        else
            a.t ← QNaN
            a.e ← -fbits(prec)
            a.f ← af
        endif
    elseif ae = 0 then
        if af = 0 then
            a.t ← ZERO
        else
            a.t ← NORM
            a.e ← 1-ebias(prec)-fbits(prec)
            a.f ← 0 11 af
        endif
    else
        a.t ← NORM
        a.e ← ae-ebias(prec)-fbits(prec)
        a.f ← 1 11 af
    endif
enddef

def a ← DEFAULTQNaN as
    a.s ← 0
    a.t ← QNaN
    a.e ← -1
    a.f ← 1
enddef

```

```

def a ← DEFAULTSNAN as
  a.s ← 0
  a.t ← SNAN
  a.e ← -1
  a.f ← 1
enddef

def fadd{a,b} as faddr{a,b,N} enddef

def c ← faddr{a,b,round} as
  if a.t=NORM and b.t=NORM then
    // d,e are a,b with exponent aligned and fraction adjusted
    if a.e > b.e then
      d ← a
      e.t ← b.t
      e.s ← b.s
      e.e ← a.e
      e.f ← b.f || 0a.e-b.e
    else if a.e < b.e then
      d.t ← a.t
      d.s ← a.s
      d.e ← b.e
      d.f ← a.f || 0b.e-a.e
      e ← b
    endif
    c.t ← d.t
    c.e ← d.e
    if d.s = e.s then
      c.s ← d.s
      c.f ← d.f + e.f
    elseif d.f > e.f then
      c.s ← d.s
      c.f ← d.f - e.f
    elseif d.f < e.f then
      c.s ← e.s
      c.f ← e.f - d.f
    else
      c.s ← r=F
      c.t ← ZERO
    endif
  // priority is given to b operand for NaN propagation
  elseif (b.t=SNAN) or (b.t=QNAN) then
    c ← b
  elseif (a.t=SNAN) or (a.t=QNAN) then
    c ← a
  elseif a.t=ZERO and b.t=ZERO then
    c.t ← ZERO
    c.s ← (a.s and b.s) or (round=F and {a.s or b.s})
  // NULL values are like zero, but do not combine with ZERO to alter sign
  elseif a.t=ZERO or a.t=NULL then
    c ← b
  elseif b.t=ZERO or b.t=NULL then
    c ← a
  elseif a.t=INFINITY and b.t=INFINITY then

```

```

if a.s ≠ b.s then
    c ← DEFAULTSNAN // Invalid
else
    c ← a
endif
elseif a.t=INFINITY then
    c ← a
elseif b.t=INFINITY then
    c ← b
else
    assert FALSE // should have covered all the cases above
endif
enddef

def b ← fneg(a) as
    b.s ← -a.s
    b.t ← a.t
    b.e ← a.e
    b.f ← a.f
enddef

def fsub(a,b) as fsubr(a,b,N) enddef

def fsubr(a,b,round) as faddr(a,fneg(b),round) enddef

def frsub(a,b) as frsubr(a,b,N) enddef

def frsubr(a,b,round) as faddr(fneg(a),b,round) enddef

def c ← fcom(a,b) as
    if (a.t=SNAN) or (a.t=QNAN) or (b.t=SNAN) or (b.t=QNAN) then
        c ← U
    elseif a.t=INFINITY and b.t=INFINITY then
        if a.s ≠ b.s then
            c ← (a.s=0) ? G: L
        else
            c ← E
        endif
    elseif a.t=INFINITY then
        c ← (a.s=0) ? G: L
    elseif b.t=INFINITY then
        c ← (b.s=0) ? G: L
    elseif a.t=NORM and b.t=NORM then
        if a.s ≠ b.s then
            c ← (a.s=0) ? G: L
        else
            if a.e > b.e then
                af ← a.f
                bf ← b.f || 0a.e-b.e
            else
                af ← a.f || 0b.e-a.e
                bf ← b.f
            endif
            if af = bf then
                c ← E
            endif
        endif
    endif
enddef

```

```

        else
            c ← ((a.s=0) ^ (af > bf)) ? G : L
        endif
    endif
elseif a.t=NORM then
    c ← {a.s=0} ? G: L
elseif b.t=NORM then
    c ← {b.s=0} ? G: L
elseif a.t=ZERO and b.t=ZERO then
    c ← E
else
    assert FALSE // should have covered all the cases above
endif
enddef

def c ← fmul(a,b) as
    if a.t=NORM and b.t=NORM then
        c.s ← a.s ^ b.s
        c.t ← NORM
        c.e ← a.e + b.e
        c.f ← a.f * b.f
    // priority is given to b operand for NaN propagation
    elseif (b.t=SNAN) or (b.t=QNAN) then
        c.s ← a.s ^ b.s
        c.t ← b.t
        c.e ← b.e
        c.f ← b.f
    elseif (a.t=SNAN) or (a.t=QNAN) then
        c.s ← a.s ^ b.s
        c.t ← a.t
        c.e ← a.e
        c.f ← a.f
    elseif a.t=ZERO and b.t=INFINITY then
        c ← DEFAULTSNAN // Invalid
    elseif a.t=INFINITY and b.t=ZERO then
        c ← DEFAULTSNAN // Invalid
    elseif a.t=ZERO or b.t=ZERO then
        c.s ← a.s ^ b.s
        c.t ← ZERO
    else
        assert FALSE // should have covered all the cases above
    endif
enddef

def c ← fdivr(a,b) as
    if a.t=NORM and b.t=NORM then
        c.s ← a.s ^ b.s
        c.t ← NORM
        c.e ← a.e - b.e + 256
        c.f ← |af || 0256 | / b.f
    // priority is given to b operand for NaN propagation
    elseif (b.t=SNAN) or (b.t=QNAN) then
        c.s ← a.s ^ b.s
        c.t ← b.t
        c.e ← b.e

```

```

c.f ← b.f
elseif (a.t=SNAN) or (a.t=QNAN) then
    c.s ← a.s ^ b.s
    c.t ← a.t
    c.e ← a.e
    c.f ← a.f
elseif a.t=ZERO and b.t=ZERO then
    c ← DEFAULTSNAN // Invalid
elseif a.t=INFINITY and b.t=INFINITY then
    c ← DEFAULTSNAN // Invalid
elseif a.t=ZERO then
    c.s ← a.s ^ b.s
    c.t ← ZERO
elseif a.t=INFINITY then
    c.s ← a.s ^ b.s
    c.t ← INFINITY
else
    assert FALSE // should have covered all the cases above
endif
enddef

def msb ← findmsb(a) as
    MAXF ← 218 // Largest possible f value after matrix multiply
    for j ← 0 to MAXF
        if aMAXF-1..j = (0MAXF-1-j || 1) then
            msb ← j
        endif
    endfor
enddef

def ai ← PackF(prec,a,round) as
    case a.t of
        NORM:
            msb ← findmsb(a.f)
            m ← msb-1-fbits(prec) // lsb for normal
            rdn ← -ebias(prec)-a.e-1-fbits(prec) // lsb if a denormal
            rb ← (rn > rdn) ? rn : rdn
            if rb ≤ 0 then
                aifr ← a.fmsb-1..0 || 0-rb
                eadj ← 0
            else
                case round of
                    C:
                        s ← 0msb-rb || (-a.s)rb
                    F:
                        s ← 0msb-rb || (a.s)rb
                    N, NONE:
                        s ← 0msb-rb || -a.frb || a.frbrb-1
                    X:
                        if a.frb-1..0 ≠ 0 then
                            raise FloatingPointArithmetic // Inexact
                        endif
                        s ← 0
                    Z:
                endcase
            endif
        endif
    endcase
enddef

```

```

        s ← 0
    endcase
    v ← (011a.fmsb..0) + (011s)
    if vmsb = 1 then
        aifr ← vmsb-1..rb
        eadj ← 0
    else
        aifr ← 0fbits(prec)
        eadj ← 1
    endif
endif
aen ← a.e + msb - 1 + eadj + ebias(prec)
if aen ≤ 0 then
    if round = NONE then
        ai ← a.s || 0ebits(prec) || aifr
    else
        raise FloatingPointArithmetic //Underflow
    endif
elseif aen ≥ 1ebits(prec) then
    if round = NONE then
        //default: round-to-nearest overflow handling
        ai ← a.s || 1ebits(prec) || 0fbits(prec)
    else
        raise FloatingPointArithmetic //Underflow
    endif
else
    ai ← a.s || aenebits(prec)-1..0 || aifr
endif
SNAN:
if round ≠ NONE then
    raise FloatingPointArithmetic //Invalid
endif
if -a.e < fbits(prec) then
    ai ← a.s || 1ebits(prec) || a.f-a.e-1..0 || 0fbits(prec)+a.e
else
    lsb ← a.f-a.e-1-fbits(prec)+1..0 ≠ 0
    ai ← a.s || 1ebits(prec) || a.f-a.e-1..a.e-1-fbits(prec)+2 || lsb
endif
QNAN:
if -a.e < fbits(prec) then
    ai ← a.s || 1ebits(prec) || a.f-a.e-1..0 || 0fbits(prec)+a.e
else
    lsb ← a.f-a.e-1-fbits(prec)+1..0 ≠ 0
    ai ← a.s || 1ebits(prec) || a.f-a.e-1..a.e-1-fbits(prec)+2 || lsb
endif
ZERO:
ai ← a.s || 0ebits(prec) || 0fbits(prec)
INFINITY:
ai ← a.s || 1ebits(prec) || 0fbits(prec)
endcase
defdef

```

```

def ai ← fsinkr(prec, a, round) as
  case a.t of
    NORM:
      msb ← findmsb(a.f)
      rb ← -a.e
      if rb ≤ 0 then
        aifr ← a.fmsb..0 || 0-rb
        aims ← msb - rb
      else
        case round of
          C, C.D:
            s ← 0msb-rb || (-ai.s)rb
          F, F.D:
            s ← 0msb-rb || (ai.s)rb
          N, NONE:
            s ← 0msb-rb || -ai.frb || ai.frb..1
        X:
          if ai.frb..1..0 ≠ 0 then
            raise FloatingPointArithmetic // Inexact
          endif
          s ← 0
        Z, Z.D:
          s ← 0
      endcase
      v ← (011a.fmsb..0) + (011s)
      if vmsb = 1 then
        aims ← msb + 1 - rb
      else
        aims ← msb - rb
      endif
      aifr ← vaims..rb
    endif
    if aims > prec then
      case round of
        C.D, F.D, NONE, Z.D:
          ai ← a.s || (-as)prec-1
        C, F, N, X, Z:
          raise FloatingPointArithmetic // Overflow
      endcase
    elseif a.s = 0 then
      ai ← aifr
    else
      ai ← -aifr
    endif
  ZERO:
    ai ← 0prec
  SNaN, QNaN:
    case round of
      C.D, F.D, NONE, Z.D:
        ai ← 0prec
      C, F, N, X, Z:
        raise FloatingPointArithmetic // Invalid
    endcase
  INFINITY:

```

```

        case round of
            C.D, F.D, NONE, Z.D:
                ai ← a.s || (~as)prec-1
            C, F, N, X, Z:
                raise FloatingPointArithmetic // Invalid
        endcase
    endcase
enddef

def c ← frecrest[a] as
    b.s ← 0
    b.t ← NORM
    b.e ← 0
    b.f ← 1
    c ← fest[fdiv(b,a)]
enddef

def c ← frsqrest[a] as
    b.s ← 0
    b.t ← NORM
    b.e ← 0
    b.f ← 1
    c ← fest[fsqr[fdiv(b,a)]]
enddef

def c ← fest[a] as
    if (a.t=NORM) then
        msb ← findmsb(a.f)
        a.e ← a.e + msb - 13
        a.f ← a.fmsb..msb-12 || 1
    else
        c ← a
    endif
enddef

def c ← fsqr[a] as
    if (a.t=NORM) and (a.s=0) then
        c.s ← 0
        c.t ← NORM
        if (a.e0 = 1) then
            c.e ← (a.e-127) / 2
            c.f ← sqrt(a.f || 0127)
        else
            c.e ← (a.e-128) / 2
            c.f ← sqrt(a.f || 0128)
        endif
    elseif (a.t=SNAN) or (a.t=QNAN) or a.t=ZERO or ((a.t=INFINITY) and (a.s=0)) then
        c ← a
    elseif ((a.t=NORM) or (a.t=INFINITY)) and (a.s=1) then
        c ← DEFAULTSNAN // Invalid
    else
        assert FALSE // should have covered all the cases above
    endif
enddef

```

Digital Signal Processing

The Zeus processor provides a set of operations that maintain the fullest possible use of 128-bit data paths when operating on lower-precision fixed-point or floating-point vector values. These operations are useful for several application areas, including digital signal processing, image processing and synthetic graphics. The basic goal of these operations is to accelerate the performance of algorithms that exhibit the following characteristics:

Low-precision arithmetic

The operands and intermediate results are fixed-point values represented in no greater than 64 bit precision. For floating-point arithmetic, operands and intermediate results are of 16, 32, or 64 bit precision.

The fixed-point arithmetic operations include add, subtract, multiply, divide, shifts, and set on compare.

The use of fixed-point arithmetic permits various forms of operation reordering that are not permitted in floating-point arithmetic. Specifically, commutativity and associativity, and distribution identities can be used to reorder operations. Compilers can evaluate operations to determine what intermediate precision is required to get the specified arithmetic result.

Zeus supports several levels of precision, as well as operations to convert between these different levels. These precision levels are always powers of two, and are explicitly specified in the operation code.

When specified, add, subtract, and shift operations may cause a fixed-point arithmetic exception to occur on resulting conditions such as signed or unsigned overflow. The fixed-point arithmetic exception may also be invoked upon a signed or unsigned comparison.

Sequential access to data

The algorithms are or can be expressed as operations on sequentially ordered items in memory. Scatter-gather memory access or sparse-matrix techniques are not required.

Where an index variable is used with a multiplier, such multipliers must be powers of two. When the index is of the form: $nx+k$, the value of n must be a power of two, and the values referenced should have k include the majority of values in the range $0..n-1$. A negative multiplier may also be used.

Vectorizable operations


The operations performed on these sequentially ordered items are identical and independent. Conditional operations are either rewritten to use Boolean variables or masking, or the compiler is permitted to convert the code into such a form.

Data-handling Operations

The characteristics of these algorithms include sequential access to data, which permit the use of the normal load and store operations to reference the data. Octlet and hexlet loads and stores reference several sequential items of data, the number depending on the operand precision.


The discussion of these operations is independent of byte ordering, though the ordering of bit fields within octlets and hexlets must be consistent with the ordering used for bytes. Specifically, if big-endian byte ordering is used for the loads and stores, the figures below should assume that index values increase from left to right, and for little-endian byte ordering, the index values increase from right to left. For this reason, the figures indicate different index values with different shades, rather than numbering.

When an index of the $nx+k$ form is used in array operands, where n is a power of 2, data memory sequentially loaded contains elements useful for separate operands. The "shuffle" instruction divides a triclet of data up into two hexlets, with alternate bit fields of the source triclet grouped together into the two results. An immediate field, h , in the instruction specifies which of the two regrouped hexlets to select for the result. For example, two X.SHUFFLE.256 rd=rc,rb,32,128,h operations rearrange the source triclet (c,b) into two hexlets as follows:




In the shuffle operation, two hexlet registers specify the source triclet, and one of the two result hexlets are specified as hexlet register.


The example above directly applies to the case where n is 2. When n is larger, shuffle operations can be used to further subdivide the sequential stream. For example, when n is 4, we need to deal out 4 sets of doublet operands, as shown in the figure below:²



When an array result of computation is accessed with an index of the form $nx+k$, for n a power of 2, the reverse of the “deal” operation needs to be performed on vectors of results to interleave them for storage in sequential order. The “shuffle” operation interleaves the bit fields of two octlets of results into a single hexlet. For example a X.SHUFFLE.16 operation combines two octlets of doublet fields into a hexlet as follows:



For larger values of n , a series of shuffle operations can be used to combine additional sets of fields, similarly to the mechanism used for the deal operations. For example, when n is 4, we need to shuffle up 4 sets of doublet operands, as shown in the figure below:³




When the index of a source array operand or a destination array result is negated, or in other words, if of the form $nx+k$ where n is negative, the elements of the array must be arranged

²An example of the use of a four-way deal can be found in the appendix: Digital Signal Processing Applications: Conversion of Color to Monochrome

³An example of the use of a four-way shuffle can be found in the appendix: Digital Signal Processing Applications: Conversion of Monochrome to Color


in reverse order. The “swizzle” operation can reverse the order of the bit fields in a hexlet. For example, a X.SWIZZLE rd=rc,127,112 operation reverses the doublets within a hexlet:



In some cases, it is desirable to use a group instruction in which one or more operands is a single value, not an array. The “swizzle” operation can also copy operands to multiple locations within a hexlet. For example, a X.SWIZZLE 15,0 operation copies the low-order 16 bits to each doublet within a hexlet.

Variations of the deal and shuffle operations are also useful for converting from one precision to another. This may be required if one operand is represented in a different precision than another operand or the result, or if computation must be performed with intermediate precision greater than that of the operands, such as when using an integer multiply.


When converting from a higher precision to a lower precision, specifically when halving the precision of a hexlet of bit fields, half of the data must be discarded, and the bit fields packed together. The “compress” operation is a variant of the “deal” operation, in which the operand is a hexlet, and the result is an octlet. An arbitrary half-sized sub-field of each bit field can be selected to appear in the result. For example, a selection of bits 19..4 of each quadlet in a hexlet is performed by the X.COMPRESS rd=rc,16,4 operation:



When converting from lower-precision to higher-precision, specifically when doubling the precision of an octlet of bit fields, one of several techniques can be used, either multiply, expand, or shuffle. Each has certain useful properties. In the discussion below, m is the precision of the source operand.

The multiply operation, described in detail below, automatically doubles the precision of the result, so multiplication by a constant vector will simultaneously double the precision of the operand and multiply by a constant that can be represented in m bits.

An operand can be doubled in precision and shifted left with the “expand” operation, which is essentially the reverse of the “compress” operation. For example the X.EXPAND rd=rc,16,4 expands from 16 bits to 32, and shifts 4 bits left:



The “shuffle” operation can double the precision of an operand and multiply it by 1 (unsigned only), 2^m or 2^m+1 , by specifying the sources of the shuffle operation to be a zeroed register and the source operand, the source operand and zero, or both to be the source operand. When multiplying by 2^m , a constant can be freely added to the source operand by specifying the constant as the right operand to the shuffle.

Arithmetic Operations

The characteristics of the algorithms that affect the arithmetic operations most directly are low-precision arithmetic, and vectorizable operations. The fixed-point arithmetic operations provided are most of the functions provided in the standard integer unit, except for those that check conditions. These functions include add, subtract, bitwise Boolean operations, shift, set on condition, and multiply, in forms that take packed sets of bit fields of a specified size as operands. The floating-point arithmetic operations provided are as complete as the scalar floating-point arithmetic set. The result is generally a packed set of bit fields of the same size as the operands, except that the fixed-point multiply function intrinsically doubles the precision of the bit field.

Conditional operations are provided only in the sense that the set on condition operations can be used to construct bit masks that can select between alternate vector expressions, using the bitwise Boolean operations. All instructions operate over the entire octlet or hexlet operands, and produce a hexlet result. The sizes of the bit fields supported are always powers of two.

Galois Field Operations

Zeus provides a general software solution to the most common operations required for Galois Field arithmetic. The instructions provided include a polynomial multiply, with the polynomial specified as one register operand. This instruction can be used to perform CRC generation and checking, Reed-Solomon code generation and checking, and spread-spectrum encoding and decoding.

Software Conventions

The following section describes software conventions that are to be employed at software module boundaries, in order to permit the combination of separately compiled code and to provide standard interfaces between application, library and system software. Register usage and procedure call conventions may be modified, simplified or optimized when a single compilation encloses procedures within a compilation unit so that the procedures have no external interfaces. For example, internal procedures may permit a greater number of register-passed parameters, or have registers allocated to avoid the need to save registers at procedure boundaries, or may use a single stack or data pointer allocation to suffice for more than one level of procedure call.

Register Usage

All Zeus registers are identical and general-purpose; there is no dedicated zero-valued register, and no dedicated floating-point registers. However, some procedure-call-oriented instructions imply usage of registers zero (0) and one (1) in a manner consistent with the conventions described below. By software convention, the non-specific general registers are used in more specific ways.

register number	assembler names	usage	how saved
0	lp, r0	link pointer	caller
1	dp, r1	data pointer	caller
2-9	r2-r9	parameters	caller
10-31	r10-r31	temporary	caller
32-61	r32-r61	saved	callee
62	fp, r62	frame pointer	callee
63	sp, r63	stack pointer	callee
register usage			

At a procedure call boundary, registers are saved either by the caller or callee procedure, which provides a mechanism for leaf procedures to avoid needing to save registers. Compilers may choose to allocate variables into caller or callee saved registers depending on how their lifetimes overlap with procedure calls.

Procedure Calling Conventions

Procedure parameters are normally allocated in registers, starting from register 2 up to register 9. These registers hold up to 8 parameters, which may each be of any size from one byte to sixteen bytes (hexlet), including floating-point and small structure parameters. Additional parameters are passed in memory, allocated on the stack. For C procedures which use varargs.h or stdarg.h and pass parameters to further procedures, the compilers must leave room in the stack memory allocation to save registers 2 through 9 into memory contiguously with the additional stack memory parameters, so that procedures such as _doprnt can refer to the parameters as an array.

Procedure return values are also allocated in registers, starting from register 2 up to register 9. Larger values are passed in memory, allocated on the stack.


There are several pointers maintained in registers for the procedure calling conventions: lp, sp, dp, fp.

The lp register contains the address to which the callee should return to at the conclusion of the procedure. If the procedure is also a caller, the lp register will need to be saved on the stack, once, before any procedure call, and restored, once, after all procedure calls. The procedure returns with a branch instruction, specifying the lp register.

The sp register is used to form addresses to save parameter and other registers, maintain local variables, i.e., data that is allocated as a LIFO stack. For procedures that require a stack, normally a single allocation is performed, which allocates space for input parameters, local variables, saved registers, and output parameters all at once. The sp register is always hexlet aligned.

The dp register is used to address pointers, literals and static variables for the procedure. The dp register points to a small (approximately 4096-entry) array of pointers, literals, and statically-allocated variables, which is used locally to the procedure. The uses of the dp register are similar to the use of the gp register on a Mips R-series processor, except that each procedure may have a different value, which expands the space addressable by small offsets from this pointer. This is an important distinction, as the offset field of Zeus load and store instructions are only 12 bits. The compiler may use additional registers and/or indirect pointers to address larger regions for a single procedure. The compiler may also share a single dp register value between procedures which are compiled as a single unit (including procedures which are externally callable), eliminating the need to save, modify and restore the dp register for calls between procedures which share the same dp register value.

Load- and store- immediate-aligned instructions, specifying the dp register as the base register, are generally used to obtain values from the dp region. These instructions shift the immediate value by the logarithm of the size of the operand, so loads and stores of large operands may reach farther from the dp register than of small operands. The size of the addressable region is maximized if the elements to be placed in the dp region are sorted according to size, with the smallest elements placed closest to the dp base. At points where the size changes, appropriate padding is added to keep elements aligned to memory boundaries matching the size of the elements. Using this technique, the maximum size of the dp region is always at least 4096 items, and may be larger when the dp area is composed of a mixture of data sizes.



Alignment within dp region

The dp register mechanism also permits code to be shared, with each static instance of the dp region assigned to a different address in memory. In conjunction with position-independent or pc-relative branches, this allows library code to be dynamically relocated and shared between processes.

To implement an inter-module (separately compiled) procedure call, the lp register is loaded with the entry point of the procedure, and the dp register is loaded with the value of the dp register required for the procedure. These two values are located adjacent to each other as a pair of octlet quantities in the dp region for the calling procedure. For a statically-linked inter-module procedure call, the linker fills in the values at link time. However, this mechanism also provides for dynamic linking, by initially filling in the lp and dp fields in the data structure to invoke the dynamic linker. The dynamic linker can use the contents of the lp and/or dp registers to determine the identity of the caller and callee, to find the location to fill in the pointers and resume execution. Specifically, the lp value is initially set to point to an entry point in the dynamic linker, and the dp value is set to point to itself: the location of the lp and dp values in the dp region of the calling procedure. The identity of the procedure

can be discovered from a string following the dp pointer, or a separate table, indexed by the dp pointer.

The fp register is used to address the stack frame when the stack size varies during execution of a procedure, such as when using the GNU C alloca function. When the stack size can be determined at compile time, the sp register is used to address the stack frame and the fp register may be used for any other general purpose as a callee-saved register.

Typical static-linked, Intra-module calling sequence:

caller (non-leaf):

```
caller: AADDI          sp@-size      // allocate caller stack frame
       S.I.64.A        lp,sp,off    // save original lp register
       ... (callee using same dp as caller)
       B.LINK.I        callee
       ...
       ... (callee using same dp as caller)
       B.LINK.I        callee
       ...
       LI.64.A          lp=sp,off    // restore original lp register
       AADDI           sp@size      // deallocate caller stack frame
       B                lp          // return
```

callee (leaf):

```
callee: ... (code using dp)
       B             lp          // return
```

Procedures that are compiled together may share a common data region, in which case there is no need to save, load, and restore the dp region in the callee, assuming that the callee does not modify the dp register. The pc-relative addressing of the B.LINK.I instruction permits the code region to be position-independent.

Minimum static-linked, Intra-module calling sequence:

caller (non-leaf):

```
caller: ACOPY          r31=lp      // save original lp register
       ... (callee using same dp as caller)
       B.LINK.I        callee
       ...
       ... (callee using same dp as caller)
       B.LINK.I        callee
       ...
       B               r31         // return
```

callee (leaf):

```
callee: ... (code using dp, r31 unused)
       B             lp          // return
```

When all the callee procedures are intra-module, the stack frame may also be eliminated from the caller procedure by using "temporary" caller save registers not utilized by the callee leaf procedures. In addition to the lp value indicated above, this usage may include other values and variables that live in the caller procedure across callee procedure calls.

Typical dynamic-linked, Inter-module calling sequence:

caller (non-leaf):

caller:	AADDI	sp@-size	// allocate caller stack frame
	S.I.64.A	lp,sp,off	// save original lp register
	S.I.64.A	dp,sp,off	// save original dp register
	... (code using dp)		
	L.I.64.A	lp=dp.off	// load lp
	L.I.64.A	dp=dp.off	// load dp
	B.LINK	lp=lp	// invoke callee procedure
	L.I.64.A	dp=sp.off	// restore dp register from stack
	... (code using dp)		
	L.I.64.A	lp=sp.off	// restore original lp register
	AADDI	sp=size	// deallocate caller stack frame
	B	lp	// return

callee (leaf):

callee:	... (code using dp)		
	B	lp	// return

The load instruction is required in the caller following the procedure call to restore the dp register. A second load instruction also restores the lp register, which may be located at any point between the last procedure call and the branch instruction which returns from the procedure.

System and Privileged Library Calls

It is an objective to make calls to system facilities and privileged libraries as similar as possible to normal procedure calls as described above. Rather than invoke system calls as an exception, which involves significant latency and complication, we prefer to use a modified procedure call in which the process privilege level is quietly raised to the required level. To provide this mechanism safely, interaction with the virtual memory system is required.


Such a procedure must not be entered from anywhere other than its legitimate entry point, to prohibit entering a procedure after the point at which security checks are performed or with invalid register contents, otherwise the access to a higher privilege level can lead to a security violation. In addition, the procedure generally must have access to memory data, for which addresses must be produced by the privileged code. To facilitate generating these addresses, the branch-gateway instruction allows the privileged code procedure to rely on the fact that a single register has been verified to contain a pointer to a valid memory region.

The branch-gateway instruction ensures both that the procedure is invoked at a proper entry point, and that other registers such as the data pointer and stack pointer can be properly set. To ensure this, the branch-gateway instruction retrieves a “gateway” directly from the protected virtual memory space. The gateway contains the virtual address of the entry point of the procedure and the target privilege level. A gateway can only exist in regions of the virtual address space designated to contain them, and can only be used to access privilege levels at or below the privilege level at which the memory region can be written to ensure that a gateway cannot be forged.

The branch-gateway instruction ensures that register 1 (dp) contains a valid pointer to the gateway for this target code address by comparing the contents of register 0 (lp) against the

gateway retrieved from memory and causing an exception trap if they do not match. By ensuring that register 1 points to the gateway, auxiliary information, such as the data pointer and stack pointer can be set by loading values located by the contents of register 1. For example, the eight bytes following the gateway may be used as a pointer to a data region for the procedure.

Before executing the branch-gateway instruction, register 1 must be set to point at the gateway, and register 0 must be set to the address of the target code address plus the desired privilege level. A "L.I.64.LA r0=r1,0" instruction is one way to set register 0, if register 1 has already been set, but any means of getting the correct value into register 0 is permissible.



Similarly, a return from a system or privileged routine involves a reduction of privilege. This need not be carefully controlled by architectural facilities, so a procedure may freely branch to a less-privileged code address. Normally, such a procedure restores the stack frame, then uses the branch-down instruction to return.

Typical dynamic-linked, Inter-gateway calling sequence:

caller:

caller:	AADDI	sp@-size	// allocate caller stack frame
	S.I.64.A	lp,sp,off	
	S.I.64.A	dp,sp,off	
	...		
	L.I.64.A	lp=dp,off	// load lp
	L.I.64.A	dp=dp,off	// load dp
	B.GATE		
	L.I.64.A	dp,sp,off	
	... (code using dp)		
	L.I.64.A	lp=sp,off	// restore original lp register
	AADDI	sp=size	// deallocate caller stack frame
	B	lp	// return

callee (non-leaf):

calee:	L.I.64.A	dp=dp,off	// load dp with data pointer
	S.I.64.A	sp,dp,off	
	L.I.64.A	sp=dp,off	// new stack pointer
	S.I.64.A	lp,sp,off	
	S.I.64.A	dp,sp,off	
	... (using dp)		
	L.I.64.A	dp,sp,off	
	... (code using dp)		
	L.I.64.A	lp=sp,off	// restore original lp register
	L.I.64.A	sp=sp,off	// restore original sp register
	B.DOWN	lp	

callee (leaf, no stack):

calee:	... (using dp)		
	B.DOWN	lp	

It can be observed that the calling sequence is identical to that of the inter-module calling sequence shown above, except for the use of the B.GATE instruction instead of a B.LINK instruction. Indeed, if a B.GATE instruction is used when the privilege level in the lp register is not higher than the current privilege level, the B.GATE instruction performs an identical function to a B.LINK.

The callee, if it uses a stack for local variable allocation, cannot necessarily trust the value of the sp passed to it, as it can be forged. Similarly, any pointers which the callee provides should not be used directly unless they are verified to point to regions which the callee should be permitted to address. This can be avoided by defining application programming interfaces (APIs) in which all values are passed and returned in registers, or by using a trusted, intermediate privilege wrapper routine to pass and return parameters. The method described below can also be used.

It can be useful to have highly privileged code call less-privileged routines. For example, a user may request that errors in a privileged routine be reported by invoking a user-supplied error-logging routine. To invoke the procedure, the privilege can be reduced via the branch-down instruction. The return from the procedure actually requires an increase in privilege, which must be carefully controlled. This is dealt with by placing the procedure call within a lower-privilege procedure wrapper, which uses the branch-gateway instruction to return to

the higher privilege region after the call through a secure re-entry point. Special care must be taken to ensure that the less-privileged routine is not permitted to gain unauthorized access by corruption of the stack or saved registers, such as by saving all registers and setting up a new stack frame (or restoring the original lower-privilege stack) that may be manipulated by the less-privileged routine. Finally, such a technique is vulnerable to an unprivileged routine attempting to use the re-entry point directly, so it may be appropriate to keep a privileged state variable which controls permission to enter at the re-entry point.

Instruction Scheduling

The next section describes detailed pipeline organization for Zeus, which has a significant influence on instruction scheduling. Here we will elaborate some general rules for effective scheduling by a compiler. Specific information on numbers of functional units, functional unit parallelism and latency is quite implementation-dependent, values indicated here are valid for Zeus's first implementation.

Separate Addressing from Execution

Zeus has separate function units to perform addressing operations (A, L, S, B instructions) from execution operations (G, X, E, W instructions). When possible, Zeus will execute all the addressing operations of an instruction stream, deferring execution of the execution operations until dependent load instructions are completed. Thus, the latency of the memory system is hidden, so long as addressing operations themselves do not need to wait for memory.

Software Pipeline

Instructions should generally be scheduled so that previous operations can be completed at the time of issue. When this is not possible, the processor inserts sufficient empty cycles to perform the instructions precisely - explicit no-operation instructions are not required.

Multiple Issue

Zeus can issue up to two addressing operations and up to two execution operations per cycle per thread. Considering functional unit parallelism, described below, as many of four instruction issues per cycle are possible per thread.

Functional Unit parallelism

Zeus has separate function units for several classes of execution operations. An A unit performs scalar add, subtract, boolean, and shift-add operations for addressing and branch calculations. The remaining functional units are execution resources, which perform operations subsequent to memory loads and which operate on values in a parallel, partitioned form. A G unit performs add, subtract, boolean, and shift-add operations. An X unit performs general shift operations. An E unit performs multiply and floating-point operations. A T unit performs table-look-up operations.

Each instruction uses one or more of these units, according to the table below.

Instruction	A	G	X	E	T
A	x				
B	x				
L	x				
S	x				
G		x			
X			x		
E			x	x	
W.TRANSLATE	x				x
W.MULMAT	x		x	x	
W.SWITCH	x		x		

Latency

The latency of each functional unit depends on what operation is performed in the unit, and where the result is used. The aggressive nature of the pipeline makes it difficult to characterize the latency of each operation with a single number. Because the addressing unit is decoupled from the execution unit, the latency of load operations is generally hidden, unless the result of a load instruction must be returned to the addressing unit. Store instructions must be able to compute the address to which the data is to be stored in the addressing unit, but the data will not be irrevocably stored until the data is available and it is valid to retire the store instruction. However, under certain conditions, data may be forwarded from a store instruction to subsequent load instructions, once the data is available.

The latency of each of these units, for the initial Zeus implementation is indicated below:

Unit	instruction	Latency rules
A	A	1 cycle
	L	Address operands must be ready to issue, 4 cycles to A unit, 0 to G, X, E, T units
	S	Address operands must be ready to issue, Store occurs when data is ready and instruction may be retired.
	B	Conditional branch operands may be provided from the A unit (64-bit values), or the G unit (128- bit values). 4 cycles for mispredicted branch
G	W	Address operand must be ready to issue,
	G	1 cycle
X	X, W.SWITCH	1 cycle for data operands, 2 cycles for shift amount or control operand
E	E, W.MULMAT	4 cycles
T	W.TRANSLATE	1 cycles


Pipeline Organization

Zeus performs all instructions *as if* executed one-by-one, in-order, with precise exceptions always available. Consequently, code that ignores the subsequent discussion of Zeus pipeline implementations will still perform correctly. However, the highest performance of the Zeus processor is achieved only by matching the ordering of instructions to the characteristics of the pipeline. In the following discussion, the general characteristics of all Zeus implementations precede discussion of specific choices for specific implementations.


Classical Pipeline Structures

Pipelining in general refers to hardware structures that overlap various stages of execution of a series of instructions so that the time required to perform the series of instructions is less than the sum of the times required to perform each of the instructions separately. Additionally, pipelines carry to connotation of a collection of hardware structures which have a simple ordering and where each structure performs a specialized function.


The diagram below shows the timing of what has become a canonical pipeline structure for a simple RISC processor, with time on the horizontal axis increasing to the right, and successive instructions on the vertical axis going downward. The stages I, R, E, M, and W refer to units which perform instruction fetch, register file fetch, execution, data memory fetch, and register file write. The stages are aligned so that the result of the execution of an instruction may be used as the source of the execution of an immediately following instruction, as seen by the fact that the end of an E stage (bold in line 1) lines up with the beginning of the E stage (bold in line 2) immediately below. Also, it can be seen that the result of a load operation executing in stages E and M (bold in line 3) is not available in the immediately following instruction (line 4), but may be used two cycles later (line 5); this is the cause of the load delay slot seen on some RISC processors.




In the diagrams below, we simplify the diagrams somewhat by eliminating the pipe stages for instruction fetch, register file fetch, and register file write, which can be understood to precede and follow the portions of the pipelines diagrammed. The diagram above is shown again in this new format, showing that the canonical pipeline has very little overlap of the actual execution of instructions.



A superscalar pipeline is one capable of simultaneously issuing two or more instructions which are independent, in that they can be executed in either order and separately, producing the same result as if they were executed serially. The diagram below shows a two-way superscalar processor, where one instruction may be a register-to-register operation (using stage E) and the other may be a register-to-register operation (using stage A) or a memory load or store (using stages A and M).



A superpipelined pipeline is one capable of issuing simple instructions frequently enough that the result of a simple instruction must be independent of the immediately following one or more instructions. The diagram below shows a two-cycle superpipelined implementation:



In the diagrams below, pipeline stages are labelled with the type of instruction that may be performed by that stage. The position of the stage further identifies the function of that stage, as for example a load operation may require several L stages to complete the instruction.


Superstring Pipeline

Zeus architecture provides for implementations designed to fetch and execute several instructions in each clock cycle. For a particular ordering of instruction types, one instruction of each type may be issued in a single clock cycle. The ordering required is A, L, E, S, B; in other words, a register-to-register address calculation, a memory load, a register-to-register data calculation, a memory store, and a branch. Because of the organization of the pipeline, each of these instructions may be serially dependent. Instructions of type E include the fixed-point execute-phase instructions as well as floating-point and digital signal processing instructions. We call this form of pipeline organization "superstring,"⁴ because of the ability to issue a string of dependent instructions in a single clock cycle, as distinguished

⁴Readers with a background in theoretical physics may have seen this term in an other, unrelated, context.

from superscalar or superpipelined organizations, which can only issue sets of independent instructions.


These instructions take from one to four cycles of latency to execute, and a branch prediction mechanism is used to keep the pipeline filled. The diagram below shows a box for the interval between issue of each instruction and the completion. Bold letters mark the critical latency paths of the instructions, that is, the periods between the required availability of the source registers and the earliest availability of the result registers. The A-L critical latency path is a special case, in which the result of the A instruction may be used as the base register of the L instruction without penalty. E instructions may require additional cycles of latency for certain operations, such as fixed-point multiply and divide, floating-point and digital signal processing operations.




Superspring Pipeline

Zeus architecture provides an additional refinement to the organization defined above, in which the time permitted by the pipeline to service load operations may be flexibly extended. Thus, the front of the pipeline, in which A, L and B type instructions are handled, is decoupled from the back of the pipeline, in which E, and S type instructions are handled. This decoupling occurs at the point at which the data cache and its backing memory is referenced; similarly, a FIFO that is filled by the instruction fetch unit decouples instruction cache references from the front of the pipeline shown above. The depth of the FIFO structures is implementation-dependent, i.e. not fixed by the architecture.

The diagram below indicates why we call this pipeline organization feature "superspring," an extension of our superstring organization.



With the super-spring organization, the latency of load instructions can be hidden, as execute instructions are deferred until the results of the load are available. Nevertheless, the execution unit still processes instructions in normal order, and provides precise exceptions.



Superthread Pipeline

This technique is not employed in the initial Zeus implementation, though it was present in an earlier prototype implementation.

A difficulty of superpipelining is that dependent operations must be separated by the latency of the pipeline, and for highly pipelined machines, the latency of simple operations can be quite significant. The Zeus "superthread" pipeline provides for very highly pipelined implementations by alternating execution of two or more independent threads. In this context, a thread is the state required to maintain an independent execution; the architectural state required is that of the register file contents, program counter, privilege level, local TB, and when required, exception status. Ensuring that only one thread may handle an exception at one time may minimize the latter state, exception status. In order to ensure that all threads make reasonable forward progress, several of the machine resources must be scheduled fairly.

An example of a resource that is critical that it be fairly shared is the data memory/cache subsystem. In a prototype implementation, Zeus is able to perform a load operation only on every second cycle, and a store operation only on every fourth cycle. Zeus schedules these fixed timing resources fairly by using a round-robin schedule for a number of threads that is relatively prime to the resource reuse rates. For this implementation, five simultaneous threads of execution ensure that resources which may be used every two or four cycles are fairly shared by allowing the instructions which use those resources to be issued only on every second or fourth issue slot for that thread.

In the diagram below, the thread number which issues an instruction is indicated on each clock cycle, and below it, a list of which functional units may be used by that instruction. The diagram repeats every 20 cycles, so cycle 20 is similar to cycle 0, cycle 21 is similar to cycle 1, etc. This schedule ensures that no resource conflict occur between threads for these resources. Thread 0 may issue an E, L, S or B on cycle 0, but on its next opportunity, cycle 5, may only issue E or B, and on cycle 10 may issue E, L or B, and on cycle 15, may issue E or B.

cycle thread	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
0	1	2	3	4	0	1	2	3	4	0	1	2	3	4	0	1	2	3	4	
E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	
L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L		
S		S		S		S		S		S		S		S		S		S		
B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	

Superthread pipeline

When seen from the perspective of an individual thread, the resource use diagram looks similar to that of the collection. Thus an individual thread may use the load unit every two instructions, and the store unit every four instructions.

cycle thread	0	5	10	15	20	25	30	35	40	45	50	55	60	65	70	75	80	85	90	95
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	
L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L	L		
S		S		S		S		S		S		S		S		S		S		
B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	

Superthread pipeline

A Zeus Superthread pipeline, with 5 simultaneous threads of execution, permits simple operations, such as register-to-register add (G.ADD), to take 5 cycles to complete; allowing for an extremely deeply pipelined implementation.

Simultaneous Multithreading

The initial Zeus implementation performs simultaneous multithreading among 4 threads. Each of the 4 threads share a common memory system, a common T unit. Pairs of threads share two G units, one X unit, and one E unit. Each thread individually has two A units. A fair allocation scheme balances access to the shared resources by the four threads.

Branch/fetch Prediction

Zeus does not have delayed branch instructions, and so relies upon branch or fetch prediction to keep the pipeline full around unconditional and conditional branch instructions. In the simplest form of branch prediction, as in Zeus's first implementation, a taken conditional backward (toward a lower address) branch predicts that a future execution of the same branch will be taken. More elaborate prediction may cache the source and target addresses of multiple branches, both conditional and unconditional, and both forward and reverse.

The hardware prediction mechanism is tuned for optimizing conditional branches that close loops or express frequent alternatives, and will generally require substantially more cycles when executing conditional branches whose outcome is not predominately taken or not-taken. For such cases of unpredictable conditional results, the use of code that avoids conditional branches in favor of the use of compare-set and multiplex instructions may result in greater performance.

Under some conditions, the above technique may not be applicable, for example if the conditional branch "guards" code which cannot be performed when the branch is taken. This may occur, for example, when a conditional branch tests for a valid (non-zero) pointer and the conditional code performs a load or store using the pointer. In these cases, the conditional branch has a small positive offset, but is unpredictable. A Zeus pipeline may handle this case as if the branch is always predicted to be not taken, with the recovery of a misprediction causing cancellation of the instructions which have already been issued but not completed which would be skipped over by the taken conditional branch. This "conditional-skip" optimization is performed by the initial Zeus implementation and requires no specific architectural feature to access or implement.

A Zeus pipeline may also perform "branch-return" optimization, in which a branch-link instruction saves a branch target address that is used to predict the target of the next returning branch instruction. This optimization may be implemented with a depth of one (only one return address kept), or as a stack of finite depth, where a branch and link pushes onto the stack, and a branch-register pops from the stack. This optimization can eliminate the misprediction cost of simple procedure calls, as the calling branch is susceptible to hardware prediction, and the returning branch is predictable by the branch-return optimization. Like the conditional-skip optimization described above, this feature is performed by the initial Zeus implementation and requires no specific architectural feature to access or implement.

Zeus implements two related instructions that can eliminate or reduce branch delays for conditional loops, conditional branches, and computed branches. The "branch-hint" instruction has no effect on architectural state, but informs the instruction fetch unit of a potential future branch instruction, giving the addresses of both the branch instruction and of the branch target. The two forms of the instruction specify the branch instruction address relative to the current address as an immediate field, and one form (branch-hint-immediate) specifies the branch target address relative to the current address as an immediate field, and the other (branch-hint) specifies the branch target address from a general register. The branch-hint-immediate instruction is generally used to give advance notice to the instruction

fetch unit of a branch-conditional instruction, so that instructions at the target of the branch can be fetched in advance of the branch-conditional instruction reaching the execution pipeline. Placing the branch hint as early as possible, and at a point where the extra instruction will not reduce the execution rate optimizes performance. In other words, an optimizing compiler should insert the branch-hint instruction as early as possible in the basic block where the parcel will contain at most one other "front-end" instruction.

Additional Load and Execute Resources

Studies of the dynamic distribution of Zeus instructions on various benchmark suites indicate that the most frequently-issued instruction classes are load instructions and execute instructions. In a high-performance Zeus implementation, it is advantageous to consider execution pipelines in which the ability to target the machine resources toward issuing load and execute instructions is increased.

One of the means to increase the ability to issue execute-class instructions is to provide the means to issue two execute instructions in a single-issue string. The execution unit actually requires several distinct resources, so by partitioning these resources, the issue capability can be increased without increasing the number of functional units, other than the increased register file read and write ports. The partitioning favored for the initial implementation places all instructions that involve shifting and shuffling in one execution unit, and all instructions that involve multiplication, including fixed-point and floating-point multiply and add in another unit. Resources used for implementing add, subtract, and bitwise logical operations may be duplicated, being modest in size compared to the shift and multiply units, or shared between the two units, as the operations have low-enough latency that two operations might be pipelined within a single issue cycle. These instructions must generally be independent, except perhaps that two simple add, subtract, or bitwise logical instructions may be performed dependently, if the resources for executing simple instructions are shared between the execution units.

One of the means to increase the ability to issue load-class instructions is to provide the means to issue two load instructions in a single-issue string. This would generally increase the resources required of the data fetch unit and the data cache, but a compensating solution is to steal the resources for the store instruction to execute the second load instruction. Thus, a single-issue string can then contain either two load instructions, or one load instruction and one store instruction, which uses the same register read ports and address computation resources as the basic 5-instruction string. This capability also may be employed to provide support for unaligned load and store instructions, where a single-issue string may contain as an alternative a single unaligned load or store instruction which uses the resources of the two load-class units in concert to accomplish the unaligned memory operation.

Result Forwarding

When temporally adjacent instructions are executed by separate resources, the results of the first instruction must generally be forwarded directly to the resource used to execute the second instruction, where the result replaces a value which may have been fetched from a register file. Such forwarding paths use significant resources. A Zeus implementation must

generally provide forwarding resources so that dependencies from earlier instructions within a string are immediately forwarded to later instructions, except between a first and second execution instruction as described above. In addition, when forwarding results from the execution units back to the data fetch unit, additional delay may be incurred.

Instruction Set

This section describes the instruction set in complete architectural detail. Operation codes are numerically defined by their position in the following operation code tables, and are referred to symbolically in the detailed instruction definitions. Entries that span more than one location in the table define the operation code identifier as the smallest value of all the locations spanned. The value of the symbol can be calculated from the sum of the legend values to the left and above the identifier.

Instructions that have great similarity and identical formats are grouped together. Starting on a new page, each category of instructions is named and introduced.

The Operation codes section lists each instruction by mnemonic that is defined on that page. A textual interpretation of each instruction is shown beside each mnemonic.

The Equivalences section lists additional instructions known to assemblers that are equivalent or special cases of base instructions, again with a textual interpretation of each instruction beside each mnemonic. Below the list, each equivalent instruction is defined, either in terms of a base instruction or another equivalent instruction. The symbol between the instruction and the definition has a particular meaning. If it is an arrow (\leftarrow or \rightarrow), it connects two mathematically equivalent operations, and the arrow direction indicates which form is preferred and produced in a reverse assembly. If the symbol is a (\Leftrightarrow), the form on the left is assembled into the form on the right solely for encoding purposes, and the form on the right is otherwise illegal in the assembler. The parameters in these definitions are formal; the names are solely for pattern-matching purposes, even though they may be suggestive of a particular meaning.

The Redundancies section lists instructions and operand values that may also be performed by other instructions in the instruction set. The symbol connecting the two forms is a (\Leftrightarrow), which indicates that the two forms are mathematically equivalent, both are legal, but the assembler does not transform one into the other.

The Selection section lists instructions and equivalences together in a tabular form that highlights the structure of the instruction mnemonics.

The Format section lists (1) the assembler format, (2) the C intrinsics format, (3) the bit-level instruction format, and (4) a definition of bit-level instruction format fields that are not a one-for-one match with named fields in the assembler format.

The Definition section gives a precise definition of each basic instruction.

The Exceptions section lists exceptions that may be caused by the execution of the instructions in this category.

Major Operation Codes

All instructions are 32 bits in size, and use the high order 8 bits to specify a major operation code.

31	24 23	0
major	other	
8	24	0

The major field is filled with a value specified by the following table:⁵

MAJOR	0	32	64	96	128	160	192	224
0	ARES	BEF16	LI16L	SI16L	GADDI	XDEPOSIT	EMULXI	WMULMATXIL
1	AADDI	BEF32	LI16B	SI16B	GADDI		EMULXIU	WMULMATXIB
2	AADDI.O	BEF64	LI16AL	SI16AL	GADDI.O		EMULXIM	WMULMATXILU
3	AADDI.U.O	BEF128	LI16AB	SI16AB	GADDI.U.O		EMULXIC	WMULMATXIBU
4		BLGF16	LI32L	SI32L		XDEPOSITU	EMULADDXI	WMULMATXIML
5	ASUBI	BLGF32	LI32B	SI32B	GSUBI		EMULADDXIU	WMULMATXIMB
6	ASUBI.O	BLGF64	LI32AL	SI32AL	GSUBI.O		EMULADDXIM	WMULMATXICL
7	ASUBI.U.O	BLGF128	LI32AB	SI32AB	GSUBI.U.O		EMULADDXIC	WMULMATXICB
8	ASETEI	BLF16	LI64L	SI64L	GSETEI	XWITHDRAW	ECONXIL	
9	ASETNEI	BLF32	LI64B	SI64B	GSETNEI		ECONXIB	
10	ASETANDEI	BLF64	LI64AL	SI64AL	GSETANDEI		ECONXIUL	
11	ASETANDNEI	BLF128	LI64AB	SI64AB	GSETANDNEI		ECONXIUB	
12	ASETLI	BGEF16	LI128L	SI128L	GSETLI	XWITHDRAWU	ECONXIML	
13	ASETGEI	BGEF32	LI128B	SI128B	GSETGEI		ECONXIMB	
14	ASETUJU	BGEF64	LI128AL	SI128AL	GSETUJU		ECONXICL	
15	ASETGEIU	BGEF128	LI128AB	SI128AB	GSETGEIU		ECONXICB	
16	AANDI	BE	LIU16L	SASI64AL	GANDI	XDEPOSITM	ESCALADDF16	WMULMATXL
17	ANANDI	BNE	LIU16B	SASI64AB	GNANDI		ESCALADDF32	WMULMATXB
18	AORI	BANDE	LIU16AL	SCSI64AL	GORI		ESCALADDF64	WMULMATGL
19	ANORI	BANDNE	LIU16AB	SCSI64AB	GNORI		ESCALADDX	WMULMATGB
20	AXORI	BL	LIU32L	SMSI64AL	GXORI	XSWIZZLE	EMULG8	
21	AMUX	BGE	LIU32B	SMSI64AB	GMUX		EMULG64	
22		BLU	LIU32AL	SMUXI64AL	GBOOLEAN		EMULX	
23		BGEU	LIU32AB	SMUXI64AB			EEXTRACT	
24	ACOPYI	BVF32	LIU64L		GCOPYI	XEXTRACT	EEXTRACTI	
25		BNVF32	LIU64B			XSELECT8	EEXTRACTIU	
26		BIF32	LIU64AL					WTABLEL
27		BNIF32	LIU64AB		G8		E.8	WTABLEB
28		BI	LI8	SI8	G16	XSHUFFLE	E.16	WSWITCHL
29		BLINKI	LIU8		G32	XSHIFTI	E.32	WSWITCHB
30		BHINTI			G64	XSHIFT	E.64	WMINORL
31	AMINOR	BMINOR	LMINOR	SMINOR	G128		E.128	WMINORB

major operation code field values

⁵Blank table entries cause the Reserved Instruction exception to occur.

Minor Operation Codes

For the major operation field values A.MINOR, B.MINOR, L.MINOR, S.MINOR, G.8, G.16, G.32, G.64, G.128, XSHIFTI, XSHIFT, E.8, E.16, E.32, E.64, E.128, W.MINOR.L and W.MINOR.B, the lowest-order six bits in the instruction specify a minor operation code:

31	24 23	other	6 5	0
major		other	minor	
8		18	6	

The minor field is filled with a value from one of the following tables:

A.MINOR	0	8	16	24	32	40	48	56
0		AAND	ASETE	ASETEF		ASHL	ASHLIADD	
1	AADD	AXOR	ASETNE	ASETLCGF				
2	AADD0	AOR	ASETANDE	ASETLF		ASHLIO		
3	AADDU0	AANDN	ASETANDNE	ASETGEF		ASHLUO		
4		AORN	ASET/LZ	ASETEFX				
5	ASUB	AXNOR	ASETGE/GEZ	ASETLGFX				
6	ASUB0	ANOR	ASETLU/GZ	ASETLFX		ASHRI		
7	ASUBU0	ANAND	ASETGEU/LEZ	ASETGEFX		ASHRIU		ACOM

minor operation code field values for A.MINOR

B.MINOR	0	8	16	24	32	40	48	56
0	B							
1	BLINK							
2	BHINT							
3	BDOWN							
4	BGATE							
5	BBACK							
6	BHALT							
7	BBARRIER							

minor operation code field values for B.MINOR

L.MINOR	0	8	16	24	32	40	48	56
0	L16L	L64L	LU16L	LU64L				
1	L16B	L64B	LU16B	LU64B				
2	L16AL	L64AL	LU16AL	LU64AL				
3	L16AB	L64AB	LU16AB	LU64AB				
4	L32L	L128L	LU32L	L8				
5	L32B	L128B	LU32B	LU8				
6	L32AL	L128AL	LU32AL					
7	L32AB	L128AB	LU32AB					

minor operation code field values for L.MINOR

S.MINOR	0	8	16	24	32	40	48	56
0	S16L	S64L	SAS64AL					
1	S16B	S64B	SAS64AB					
2	S16AL	S64AL	SCS64AL	SDCS64AL				
3	S16AB	S64AB	SCS64AB	SDCS64AB				
4	S32L	S128L	SMS64AL	S8				
5	S32B	S128B	SMS64AB					
6	S32AL	S128AL	SMUX64AL					
7	S32AB	S128AB	SMUX64AB					

minor operation code field values for S.MINOR

G.size	0	8	16	24	32	40	48	56
0			GSETE	GSETEF	GADDHN	GSUBHN		
1	GADD		GSETNE	GSETLGF	GADDHZ	GSUBHZ		
2	GADDO		GSETANDE	GSETLF	GADDHF	GSUBHF		
3	GADDUO		GSETANDNE	GSETGEF	GADDHC	GSUBHC		
4			GSETL/L2	GSETEFX	GADDHUN	GSUBHUN		
5	GSUB		GSETGE/GEZ	GSETLGFX	GADDHUZ	GSUBHuz		
6	GSUBO		GSETL/L/GZ	GSETLFX	GADDHUF	GSUBHUF		
7	GSUBUO		GSETGEU/LEZ	GSETGEFX	GADDHUC	GSUBHUC		

minor operation code field values for G.size

XSHIFTI	0	8	16	24	32	40	48	56
0	XSHLI	XSHLIO		XSHRI		XEXPANDI		XCOMPRESSI
1								
2								
3								
4	XSHLMI	XSHLIOU	XSHRMI	XSHRIU	XROTU	XEXPANDIU	XROTRI	XCOMPRESSIU
5								
6								
7								

minor operation code field values for XSHIFTI

XSHIFT	0	8	16	24	32	40	48	56
0	XSHL	XSHLO		XSHR		XEXPAND		XCOMPRESS
1								
2								
3								
4	XSHLM	XSHLOU	XSHRM	XSHRU	XROTL	XEXPANDU	XROTR	XCOMPRESSU
5								
6								
7								

minor operation code field values for XSHIFT

E.size	0	8	16	24	32	40	48	56
0	EMULFN	EMULADDFN	EADDFN	ESUBFN	EMUL	EMULADD	EDIVFN	ECON
1	EMULFZ	EMULADDZF	EADDZF	ESUBZF	EMULU	EMULADDU	EDIVFZ	ECONU
2	EMULFF	EMULADDFF	EADDFF	ESUBFF	EMULM	EMULADDM	EDIVFF	ECONM
3	EMULFC	EMULADDFC	EADDFC	ESUBFC	EMULC	EMULADDCC	EDIVFC	ECONC
4	EMULFX	EMULADDFX	EADDFX	ESUBFX	EMULSUM	EMULSUB	EDIVFX	EDIV
5	EMULF	EMULADDF	EADDF	ESUBF	EMULSUMU	EMULSUBU	EDIVF	EDIVU
6	EMULCF	EMULADDCC	ECONFL	ECONFL	EMULSUMM	EMULSUBM	EMULSUMF	EMULP
7	EMULSUMCF	EMULSUBCF	ECONFB	ECONFB	EMULSUMC	EMULSUBC	EMULSUBF	EUNARY

minor operation code field values for E.size

W.MINOR.order	0	8	16	24	32	40	48	56
0	WMULMAT8	WMULMATM8						
1	WMULMAT16	WMULMATM16	WMULMATF16					
2	WMULMAT32	WMULMATM32	WMULMATF32					
3	WMULMAT64	WMULMATM64	WMULMATF64					
4	WMULMAT8	WMULMATC8		WMULMATP8				
5	WMULMAT16	WMULMATC16	WMULMATCF16	WMULMATP16				
6	WMULMAT32	WMULMATC32	WMULMATCF32	WMULMATP32				
7	WMULMAT64	WMULMATC64	WMULMATCF64	WMULMATP64				

minor operation code field values for W.MINOR.L or W.MINOR.B

For the major operation field values E.MUL.X.I, E.MUL.X.I.U, E.MUL.X.I.M, E.MUL.X.I.C, E.MUL.ADD.X.I, E.MUL.ADD.X.I.U, E.MUL.ADD.X.I.M, E.MUL.ADD.X.I.C, E.CON.X.I.I, E.CON.X.I.B, E.CON.X.I.U.L, E.CON.X.I.U.B, E.CON.X.I.M.L, E.CON.X.I.M.B, E.CON.X.I.C.L, E.CON.X.I.C.B, E.EXTRACT.I, E.EXTRACT.I.U, W.MUL.MAT.X.I.U.L, W.MUL.MAT.X.I.U.B, W.MUL.MAT.X.I.M.L, W.MUL.MAT.X.I.M.B, W.MUL.MAT.X.I.C.L, and W.MUL.MAT.X.I.C.B, another six bits in the instruction specify a minor operation code, which indicates operand size, rounding, and shift amount:

31	24 23	18	17	16	15	0
major	other					minor
8		18			6	0

The minor field is filled with a value from the following table: Note that the shift amount field value shown below is the "sh" value, which is encoded in an instruction-dependent manner from the immediate field in the assembler format.

XI	0	8	16	24	32	40	48	56
0	8.F.0	8.N.0	16.F.0	16.N.0	32.F.0	32.N.0	64.F.0	64.N.0
1	8.F.1	8.N.1	16.F.1	16.N.1	32.F.1	32.N.1	64.F.1	64.N.1
2	8.F.2	8.N.2	16.F.2	16.N.2	32.F.2	32.N.2	64.F.2	64.N.2
3	8.F.3	8.N.3	16.F.3	16.N.3	32.F.3	32.N.3	64.F.3	64.N.3
4	8.Z.0	8.C.0	16.Z.0	16.C.0	32.Z.0	32.C.0	64.Z.0	64.C.0
5	8.Z.1	8.C.1	16.Z.1	16.C.1	32.Z.1	32.C.1	64.Z.1	64.C.1
6	8.Z.2	8.C.2	16.Z.2	16.C.2	32.Z.2	32.C.2	64.Z.2	64.C.2
7	8.Z.3	8.C.3	16.Z.3	16.C.3	32.Z.3	32.C.3	64.Z.3	64.C.3

minor operation code field values for EMULXI, EMULXIU, EMULXIM, EMULXIC, EMULADDXI, EMULADDXIU, EMULADDXIM, EMULADDXIC, ECONXIL, ECONXIB, ECONXIUL, ECONXIUB, ECONXIML, ECONXIMB, ECONXICL, ECONXICB, EEXTRACTI, EEXTRACTIU, WMULMATXIUL, WMULMATXIUB, WMULMATXIML, WMULMATXIMB, WMULMATXICL, and WMULMATXICB,

For the major operation field values GCOPYI, two bits in the instruction specify an operand size:

31	24 23	18	17	16	15
op	rd	sz		Imm	
8	6	2		16	

For the major operation field values G.AND.I, G.NAND.I, G.NOR.I, G.OR.I, G.XOR.I, G.ADD.I, G.ADD.I.O, G.ADD.I.UO, G.SET.AND.E.I, G.SET..AND.NE.I, G.SET.E.I, G.SET.GE.I, G.SET.L.I, G.SET.NE.I, G.SET.GE.I.U, G.SET.L.I.U, G.SUB.I, G.SUB.I.O, G.SUB.I.UO, two bits in the instruction specify an operand size:

31	24 23	18 17	12 11 10 9	
op	rd	rc	sz	Imm
8	6	6	2	10

The sz field is filled with a value from the following table:

sz	size
0	16
1	32
2	64
3	128

operand size field values for G.COPY.I, G.AND.I, G.NAND.I, G.NOR.I, G.OR.I, G.XOR.I, G.ADD.I, G.ADD.I.O, G.ADD.I.UO, G.SET.AND.E.I, G.SET..AND.NE.I, G.SET.E.I, G.SET.GE.I, G.SET.L.I, G.SET.NE.I, G.SET.GE.I.U, G.SET.L.I.U, G.SUB.I, G.SUB.I.O, G.SUB.I.UO

For the major operation field values E.8, E.16, E.32, E.64, E.128, with minor operation field value E.UNARY, another six bits in the instruction specify a unary operation code:

31	24 23	18 17	12 11	6 5	0
major	rd	rc	unary	minor	
8	6	6	6	6	6

The unary field is filled with a value from the following table:

E.UNARY	0	8	16	24	32	40	48	56
0	ESORFN	ESUMFN	ESINKFN	EFLOATFN	EDEFLATEFN	ESUM		
1	ESORFZ	ESUMFZ	ESINKFZ	EFLOATFZ	EDEFLATEFZ	ESUMU	ESINKFZD	
2	ESORFF	ESUMFF	ESINKFF	EFLOATFF	EDEFLATEFF	ELOGMOST	ESINKFFD	
3	ESORFC	ESUMFC	ESINKFC	EFLOATFC	EDEFLATEFC	ELOGMOSTU	ESINKFCU	
4	ESORFX	ESUMFX	ESINKFX	EFLOATFX	EDEFLATEFX			
5	ESORF	ESUMF	ESINKF	EFLOATF	EDEFLATEF			
6	ERSORESTFX	ERECESTFX	EABSFX	ENEGFX	EINFLATEFX		ECOPYFX	
7	ERSORESTF	ERECESTF	EABSF	ENEGF	EINFLATEF		ECOPYF	

unary operation code field values for E.UNARY.size

For the major operation field values A.MINOR and G.MINOR, with minor operation field values A.COM and G.COM, another six bits in the instruction specify a comparison operation code:

31	24 23	18 17	12 11	6 5	0
major	rd	rc	compare	minor	
8	6	6	6	6	6

The compare field is filled with a value from the following table:

x.COM	0	8	16	24	32	40	48	56
0	xCOME	xCOMEF						
1	xCOMINE	xCOMLGF						
2	xCOMANDE	xCOMLF						
3	xCOMANDNE	xCOMGEF						
4	xCOML	xCOMEFX						
5	xCOMGE	xCOMLGF						
6	xCOMLU	xCOMLFX						
7	xCOMGEU	xCOMGEFX						

compare operation code field values for A.COM.op and G.COM.op.size

General Forms

The general forms of the instructions coded by a major operation code are one of the following:

31	24 23				0
	major		offset		
8			24		
31	24 23	18 17			0
	major	rd		offset	
8		6		18	
31	24 23	18 17	12 11		0
	major	rd	rc	offset	
8		6	6	12	
31	24 23	18 17	12 11	6 5	0
	major	rd	rc	rb	ra
8		6	6	6	6

The general forms of the instructions coded by major and minor operation codes are one of the following:

31	24 23	18 17	12 11	6 5	0
	major	rd	rc	rb	minor
8		6	6	6	6
31	24 23	18 17	12 11	6 5	0
	major	rd	rc	simm	minor
8		6	6	6	6

The general form of the instructions coded by major, minor, and unary operation codes is the following:

31	24 23	18 17	12 11	6 5	0
	major	rd	rc	unary	minor
8		6	6	6	6

Register rd is either a source register or destination register, or both. Registers rc and rb are always source registers. Register ra is always a destination register.

Instruction Fetch

Definition

```

def Thread[th] as
    forever do
        catch exception
            if !(EventRegister & EventMask[th]) ≠ 0 then
                if ExceptionState=0 then
                    raise EventInterrupt
                endif
            endif
            inst ← LoadMemoryX(ProgramCounter,ProgramCounter,32,L)
            Instruction[inst]
        endcatch
        case exception of
            EventInterrupt,
            ReservedInstruction,
            AccessDisallowedByVirtualAddress,
            AccessDisallowedByTag,
            AccessDisallowedByGlobalTB,
            AccessDisallowedByLocalTB,
            AccessDetailRequiredByTag,
            AccessDetailRequiredByGlobalTB,
            AccessDetailRequiredByLocalTB,
            MissInGlobalTB,
            MissInLocalTB,
            FixedPointArithmetic,
            FloatingPointArithmetic,
            GatewayDisallowed:
                case ExceptionState of
                    0:
                        PerformException(exception)
                    1:
                        PerformException(SecondException)
                    2:
                        PerformMachineCheck(ThirdException)
                endcase
            TakenBranch:
                ContinuationState ← (ExceptionState=0) ? 0 : ContinuationState
            TakenBranchContinue:
                /* nothing */
            none, others:
                ProgramCounter ← ProgramCounter + 4
                ContinuationState ← (ExceptionState=0) ? 0 : ContinuationState
        endcase
    endforever
enddef

```

Perform Exception

Definition

```

def PerformException(exception) as
    v ← {exception > 7} ? 7 : exception
    t ← LoadMemory[ExceptionBase,ExceptionBase+Thread*128+64+8*v,64,L]
    if ExceptionState = 0 then
        u ← RegRead(3,128) || RegRead(2,128) || RegRead(1,128) || RegRead(0,128)
        StoreMemory[ExceptionBase,ExceptionBase+Thread*128,512,L,u]
        RegWrite[0,64,ProgramCounter63..2 || PrivilegeLevel
        RegWrite[1,64,ExceptionBase+Thread*128]
        RegWrite[2,64,exception]
        RegWrite[3,64,FailingAddress]
    endif
    PrivilegeLevel ← t1..0
    ProgramCounter ← t63..2 || 02
    case exception of
        AccessDetailRequiredByTag,
        AccessDetailRequiredByGlobalTB,
        AccessDetailRequiredByLocalTB:
            ContinuationState ← ContinuationState + 1
        others:
            /* nothing */
    endcase
    ExceptionState ← ExceptionState + 1
enddef

```

Instruction Decode

```

def Instruction(inst) as
    major ← inst31..24
    rd ← inst23..18
    rc ← inst17..12
    simm ← rb ← inst11..6
    minor ← ra ← inst5..0
    case major of
        ARES:
            AlwaysReserved
        A.MINOR:
            minor ← inst5..0
            case minor of
                AADD, AADD.O, AADD.OU, AAND, AANDN, ANAND, ANOR,
                AOR, AORN, AXNOR, AXOR:
                    Address(minor,rd,rc,rb)
        A.COM:
            compare ← inst11..6
            case compare of
                A.COM.E, A.COM.NE, A.COM.AND.E, A.COM.AND.NE,
                A.COM.L, A.COM.GE, A.COM.L.U, A.COM.GE.U:
                    AddressCompare(compare,rd,rc)
            others:
                raise ReservedInstruction

```

```

        endcase
ASUB, ASUB.O, ASUB.U.O,
ASET.AND.E, ASET.AND.NE, ASET.E, ASET.NE,
ASET.L, ASET.GE, ASET.LU, ASET.GE.U,
        AddressReversed(minor,rd,rc,rb)
ASHLI.ADD, ASHLI.ADD+3:
        AddressShiftLeftImmediateAdd(inst1..0,rd,rc,rb)
ASHLI.SUB, ASHLI.SUB+3:
        AddressShiftLeftImmediateSubtract(inst1..0,rd,rc,rb)
ASHLI, ASHLI.O, ASHLI.U.O, ASHRI.I, ASHRI.U, AROTR.I:
        AddressShiftImmediate(minor,rd,rc,simm)
others:
        raise ReservedInstruction
endcase
ACOPY.I
        AddressCopyImmediate(major,rd,inst17..0)
AADD.I, AADD.I.O, AADD.I.U.O, AAND.I, AOR.I, ANAND.I, ANOR.I, AXOR.I:
        AddressImmediate(major,rd,rc,inst11..0)
ASET.AND.E.I, ASET.AND.NE.I, ASET.E.I, ASET.NE.I,
ASET.L.I, ESET.GE.I, ASET.LU.I, ASET.GE.U.I,
ASUB.I, ASUB.I.O, ASUB.I.U.O:
        AddressImmediateReversed(major,rd,rc,inst11..0)
AMUX:
        AddressTernary(major,rd,rc,rb,ra)
B.MINOR:
        case minor of
        B:
            Branch(rd,rc,rb)
        B.BACK:
            BranchBack(rd,rc,rb)
        B.BARRIER:
            BranchBarrier(rd,rc,rb)
        B.DOWN:
            BranchDown(rd,rc,rb)
        B.GATE:
            BranchGateway(rd,rc,rb)
        B.HALT:
            BranchHalt(rd,rc,rb)
        B.HINT:
            BranchHint(rd,inst17..12,simm)
        B.LINK:
            BranchLink(rd,rc,rb)
        others:
            raise ReservedInstruction
endcase
BE, BNE, BL, BGE, BLU, BGE.U, BAND.E, BAND.NE:
        BranchConditional(major,rd,rc,inst11..0)
BHINTI:
        BranchHintImmediate(inst23..18,inst17..12,inst11..0)
BI:
        BranchImmediate(inst23..0)
BLINKI:
        BranchImmediateLink(inst23..0)
BEF16, BLGF16, BLF16, BGEF16,
BEF32, BLGF32, BLF32, BGEF32,
BEF64, BLGF64, BLF64, BGEF64,
BEF128, BLGF128, BLF128, BGEF128:

```

```

BranchConditionalFloatingPoint(major,rd,rc,inst11..0)
BIF32, BNIF32, BNVF32, BVF32:
BranchConditionalVisibilityFloatingPoint(major,rd,rc,inst11..0)

L.MINOR
case minor of
  L16L, LU16L, L32L, LU32L, L64L, LU64L, L128L, L8, LU8,
  L16AL, LU16AL, L32AL, LU32AL, L64AL, LU64AL, L128AL,
  L16B, LU16B, L32B, LU32B, L64B, LU64B, L128B,
  L16AB, LU16AB, L32AB, LU32AB, L64AB, LU64AB, L128AB:
    Load(minor,rd,rc,rb)
  others:
    raise ReservedInstruction
endcase

LI16L LIU16L, LI32L, LIU32L, LI64L, LIU64L, LI128L, LI8, LIU8,
LI16AL, LIU16AL, LI32AL, LIU32AL, LI64AL, LIU64AL, LI128AL,
LI16B, LIU16B, LI32B, LIU32B, LI64B, LIU64B, LI128B,
LI16AB, LIU16AB, LI32AB, LIU32AB, LI64AB, LIU64AB, LI128AB:
  LoadImmediate(major,rd,rc,inst11..0)

S.MINOR
case minor of
  S16L, S32L, S64L, S128L, S8,
  S16AL, S32AL, S64AL, S128AL,
  SAS64AL, SCS64AL, SMS64AL, SM64AL,
  S16B, S32B, S64B, S128B,
  S16AB, S32AB, S64AB, S128AB,
  SAS64AB, SCS64AB, SMS64AB, SM64AB:
    Store(minor,rd,rc,rb)
  SDCS64AB, SDCS64AL:
    StoreDoubleCompareSwap(minor,rd,rc,rb)
  others:
    raise ReservedInstruction
endcase

SI16L, SI32L, SI64L, SI128L, SI8,
SI16AL, SI32AL, SI64AL, SI128AL,
SASI64AL, SCSI64AL, SMSI64AL, SMUXI64AL,
SI16B, SI32B, SI64B, SI128B,
SI16AB, SI32AB, SI64AB, SI128AB:
SASI64AB, SCSI64AB, SMSI64AB, SMUXI64AB:
  StoreImmediate(major,rd,rc,inst11..0)

G.8, G.16, G.32, G.64, G.128:
  minor ← inst5..0
  size ← 0 11 1 11 03+major-G.8
  case minor of
    GADD, GADD.L, GADD.LU, GADD.O, GADD.OU:
      Group(minor,size,rd,rc,rb)
    GADDHC, GADDHF, GADDHN, GADDHZ,
    GADDHUC, GADDHUF, GADDHUN, GADDHUZ:
      GroupAddHalve(minor,inst1..0,size,rd,rc,rb)
    GAAA, GASA:
      GroupInplace(minor,size,rd,rc,rb)
    G.SET.AND.E, G.SET.AND.NE, G.SET.E, G.SET.NE,
    G.SET.L, G.SET.GE, G.SET.L.U, G.SET.GE.U:
    G.SUB, G.SUB.L, G.SUB.LU, G.SUB.O, G.SUB.U.O:
      GroupReversed(minor,size,ra,rb,rc)
    G.SET.E.F, G.SET.LG.F, G.SET.GE.F, G.SET.L.F,
    G.SET.E.F.X, G.SET.LG.F.X, G.SET.GE.F.X, G.SET.L.F.X:
      GroupReversedFloatingPoint(minor.op.size,

```

```

        minor.round, rd, rc, rb)
G.SHL.I.ADD..G.SHL.I.ADD+3,
    GroupShiftLeftImmediateAdd(inst1..0.size,rd,rc,rb)
G.SHL.I.SUB..G.SHL.I.SUB+3,
    GroupShiftLeftImmediateSubtract(inst1..0.size,rd,rc,rb)
G.SUBHC, G.SUBHF, G.SUBHN, G.SUBHZ,
G.SUBHUC, G.SUBHUF, G.SUBHUN, G.SUBHUZ:
    GroupSubtractHalve(minor,inst1..0.size,rd,rc,rb)
G.COM,
    compare ← inst11..6
    case compare of
        G.COM.E, G.COM.NE, G.COM.AND.E, G.COM.AND.NE,
        G.COM.L, G.COM.GE, G.COM.LU, G.COM.GE.U:
            GroupCompare(compare,size,ra,rb)
        others:
            raise ReservedInstruction
    endcase
    others:
        raise ReservedInstruction
    endcase
G.BOOLEAN..G.BOOLEAN+1:
    GroupBoolean(major,rd,rc,rb,minor)
G.COPY.I...G.COPY.I+1:
    size ← 0 11 1 11 04+inst17..16
    GroupCopyImmediate(major,size,rd,inst15..0)
G.AND.I, G.NAND.I, G.NOR.I, G.OR.I, GXOR.I,
G.ADD.I, G.ADD.I.O, G.ADD.I.U.O:
    size ← 0 11 1 11 04+inst11..10
    GroupImmediate(major,size,rd,rc,inst9..0)
G.SET.AND.E.I, G.SET.AND.NE.I, G.SET.E.I, G.SET.GE.I, G.SET.L.I,
G.SET.NE.I, G.SET.GE.I.U, G.SET.L.I.U, G.SUB.I, G.SUB.I.O, G.SUB.I.U.O:
    size ← 0 11 1 11 04+inst11..10
    GroupImmediateReversed(major,size,rd,rc,inst9..0)
G.MUX:
    GroupTernary(major,rd,rc,rb,ra)
X.SHIFT:
    minor ← inst5..2 11 02
    size ← 0 11 1 11 0(inst24 11 inst1..0)
    case minor of
        X.EXPAND, X.UEXPAND, X.SHL, X.SHL.O, X.SHL.U.O,
        X.ROTR, X.SHR, X.SHR.U,
        Crossbar(minor,size,rd,rc,rb)
        X.SHLM, X.SHR.M:
            CrossbarInplace(minor,size,rd,rc,rb)
    others:
        raise ReservedInstruction
    endcase
X.EXTRACT:
    CrossbarExtract(major,rd,rc,rb,ra)
X.DEPOSIT, X.DEPOSIT.U X.WITHDRAW X.WITHDRAW.U
    CrossbarField(major,rd,rc,inst11..6,inst5..0)
X.DEPOSIT.M:
    CrossbarFieldInplace(major,rd,rc,inst11..6,inst5..0)
X.SHIFT.I:
    minor ← inst5..0

```

```

case minor5..2 11 02 of
  X.COMPRESS.I, X.EXPAND.I, X.ROTR.I, X.SHL.I, X.SHL.I.O, X.SHL.I.U.O,
  X.SHR.I, X.COMPRESS.I.U, X.EXPAND.I.U, X.SHR.U.I:
    CrossbarShortImmediate(minor,rd,rc,simm)
  X.SHL.M.I, X.SHR.M.I:
    CrossbarShortImmediateInplace(minor,rd,rc,simm)
  others:
    raise ReservedInstruction
endcase
X.SHUFFLE.X.SHUFFLE+1:
  CrossbarShuffle(major,rd,rc,rb,simm)
X.SWIZZLE.X.SWIZZLE+3:
  CrossbarSwizzle(major,rd,rc, inst11..6,inst5..0)
X.SELECT.8:
  CrossbarTernary(major,rd,rc,rb,ra)
E.8, E.16, E.32, E.64, E.128:
  minor ← inst5..0
  size ← 0 11 1 11 03+major-E.8
  case minor of
    E.CON., E.CON.U, E.CON.M, E.CON.C,
    E.MUL., E.MUL.U, E.MUL.M, E.MUL.C,
    E.MUL.SUM, E.MUL.SUM.U, E.MUL.SUM.M, E.MUL.SUM.C,
    E.DIV, E.DIV.U, E.MUL.P:
      Ensemble{minor,size,ra,rb,rc}
    E.CON.F.L E.CON.F.B, E.CON.C.F.L, E.CON.C.F.B:
      EnsembleConvolveFloatingPoint(minor.size,rd,rc,rb)
    E.ADD.F.N, E.MUL.C.F.N, E.MUL.F.N, E.DIV.F.N,
    E.ADD.F.Z, E.MUL.C.F.Z, E.MUL.F.Z, E.DIV.F.Z,
    E.ADD.F.F, E.MUL.C.F.F, E.MUL.F.F, E.DIV.F.F,
    E.ADD.F.C, E.MUL.C.F.C, E.MUL.F.C, E.DIV.F.C,
    E.ADD.F, E.MUL.C.F, E.MUL.F, E.DIV.F,
    E.ADD.F.X, E.MUL.C.F.X, E.MUL.F.X, E.DIV.F.X,
      EnsembleFloatingPoint(minor.op, major.size, minor.round, rd, rc, rb)
    E.MULADD, E.MULADD.U, E.MULADD.M, E.MULADD.C:
      EnsembleInplace(minor.size,rd,rc,rb)
    E.MULSUB, E.MULSUB.U, E.MULSUB.M, E.MULSUB.C:
      EnsembleInplaceReversed(minor.size,rd,rc,rb)
    E.MULSUB.F, E.MULSUB.C.F:
      EnsembleInplaceReversedFloatingPoint(minor.size,rd,rc,rb)
    E.SUB.F.N, E.SUB.F.Z, E.SUB.F.F, E.SUB.F.C, E.SUB.F, E.SUB.F.X:
      EnsembleReversedFloatingPoint(minor.op, major.size,
        minor.round, rd, rc, rb)
  E.UNARY:
    case unary of
      E.SUM, E.SUM.U, E.LOG.MOST, E.LOG.MOST.U:
        EnsembleUnary(unary,rd,rc)
      E.ABS.F, E.ABS.F.X, E.COPY.F, E.COPY.F.X,
      E.DEFLATE.F, E.DEFLATE.F.N, E.DEFLATE.F.Z,
      E.DEFLATE.F.F, E.DEFLATE.F.C, E.DEFLATE.F.X:
      E.FLOAT.F, E.FLOAT.F.N, E.FLOAT.F.Z,
      E.FLOAT.F.F, E.FLOAT.F.C, E.FLOAT.F.X:
      E.INFLATE.F, E.INFLATE.F.X, E.NEG.F, E.NEG.F.X,
      E.RECEST.F, E.RECEST.F.X, E.RSQREST.F, E.RSQREST.F.X,
      E.SQR.F, E.SQR.F.N, E.SQR.F.Z, E.SQR.F.F, E.SQR.F.C, E.SQR.F.X:
      E.SUM.F, E.SUM.F.N, E.SUM.F.Z,
      E.SUM.F.F, E.SUM.F.C, E.SUM.F.X:
      E.SINK.F, E.SINK.F.Z.D, E.SINK.F.F.D, E.SINK.F.C.D, E.SINK.F.X.D,
      E.SINK.F.N, E.SINK.F.Z, E.SINK.F.F, E.SINK.F.C, E.SINK.F.X:

```

```

        EnsembleUnaryFloatingPoint(unary.op, major.size,
                                     unary.round, rd, rc)
    others:
        raise ReservedInstruction
    endcase
    others:
        raise ReservedInstruction
    endcase
E.CON.X.IL, E.CON.X.IB, E.CON.X.IUL, E.CON.X.IUB,
E.CON.X.ILM, E.CON.X.IMB, E.CON.X.ICL, E.CON.X.ICB:
    size ← 1 11 03+inst5..4
    EnsembleConvolveExtractImmediate(major,inst3..2.size,rd,rc,rb,inst1..0)
E.MULX, E.EXTRACT, E.SCALADDX:
    EnsembleExtract(major,rd,rc,rb,ra)
E.EXTRACTI, E.EXTRACTIU E.MULXI, E.MULXIU, E.MULXIM, E.MULXIC:
    size ← 1 11 03+inst5..4
    EnsembleExtractImmediate(major,inst3..2.size,rd,rc,rb,inst1..0)
E.MUL.ADD.X.I, E.MUL.ADDX.I.U, E.MUL.ADDX.I.M, E.MUL.ADDX.I.C:
    size ← 1 11 03+inst5..4
    EnsembleExtractImmediateInplace(major,inst3..2.size,rd,rc,rb,inst1..0)
E.MUL.GAL.8, E.MUL.GAL.64:
    size ← 1 11 0 3+inst26..24
    EnsembleTernary(major,size,rd,rc,rb,ra)
E.MUL.ADD.F16, E.MUL.ADD.F32, E.MUL.ADD.F64, E.MUL.ADD.F128
E.MULSUB.F16, E.MULSUB.F32, E.MULSUB.F64, E.MULSUB.F128,
E.SCALADD.F16, E.SCALADD.F32, E.SCALADD.F64:
    EnsembleTernaryFloatingPoint(major,rd,rc,rb,ra)
W.MINOR.B, W.MINOR.L:
    case minor of
        W.TRANSLATE.8, W.TRANSLATE.16, W.TRANSLATE.32, W.TRANSLATE.64:
            size ← 1 11 03+inst5..4
            WideTranslate(major,size,rd,rc,rb)
        W.MUL.MAT.8, W.MUL.MAT.16, W.MUL.MAT.32, W.MUL.MAT.64,
        W.MUL.MAT.U.8, W.MUL.MAT.U.16, W.MUL.MAT.U.32, W.MUL.MAT.U.64,
        W.MUL.MAT.M.8, W.MUL.MAT.M.16, W.MUL.MAT.M.32, W.MUL.MAT.M.64,
        W.MUL.MAT.C.8, W.MUL.MAT.C.16, W.MUL.MAT.C.32, W.MUL.MAT.C.64,
        W.MUL.MAT.P.8, W.MUL.MAT.P.16, W.MUL.MAT.P.32, W.MUL.MAT.P.64:
            size ← 1 11 03+inst5..4
            WideMultiply(major,minor,size,rd,rc,rb)
        W.MUL.MAT.F16, W.MUL.MAT.F32, W.MUL.MAT.F64,
        W.MUL.MAT.C.F16, W.MUL.MAT.C.F32, W.MUL.MAT.C.F64:
            size ← 1 11 03+inst5..4
            WideFloatingPointMultiply(major,minor,size,rd,rc,rb)
    others:
    endcase
W.MUL.MAT.X.B, W.MUL.MAT.X.L:
    WideExtract(major,ra,rb,rc,rd)
W.MUL.MAT.X.I.B, W.MUL.MAT.X.I.L, W.MUL.MAT.X.I.U.B, W.MUL.MAT.X.I.U.L,
W.MUL.MAT.X.I.M.B, W.MUL.MAT.X.I.M.L, W.MUL.MAT.X.I.C.B, W.MUL.MAT.X.I.C.L:
    size ← 1 11 03+inst5..4
    WideExtractImmediate(major,inst3..2.size,ra,rb,rc,inst1..0)
W.MUL.MAT.G.B, W.MUL.MAT.G.L:
    WideMultiplyGalois(major,rd,rc,rb,ra)
W.SWITCH.B, W.SWITCH.L:
    WideSwitch(major,rd,rc,rb,ra)
others:

```

```
raise ReservedInstruction
endcase
enddef
```

Always Reserved

This operation generates a reserved instruction exception.


Operation code

A.RES	Always reserved
-------	-----------------

Format

A.RES imm

ares(imm)



Description

The reserved instruction exception is raised. Software may depend upon this major operation code raising the reserved instruction exception in all implementations. The choice of operation code intentionally ensures that a branch to a zeroed memory area will raise an exception.

Definition

```
def AlwaysReserved as
    raise ReservedInstruction
endef
```

Exceptions

Reserved Instruction

Address

These operations perform calculations with two general register values, placing the result in a general register.

Operation codes

AADD	Address add
AADD.O	Address add signed check overflow
AADD.U.O	Address add unsigned check overflow
AAND	Address and
AANDN	Address and not
ANAND	Address not and
ANOR	Address not or
AOR	Address or
AORN	Address or not
AXNOR	Address exclusive nor
AXOR	Address xor

Redundancies

AOR rd=rc,rc	\Leftrightarrow ACOPY rd=rc
AAND rd=rc,rc	\Leftrightarrow ACOPY rd=rc
ANAND rd=rc,rc	\Leftrightarrow ANOT rd=rc
ANOR rd=rc,rc	\Leftrightarrow ANOT rd=rc
AXNOR rd=rc,rc	\Leftrightarrow ASET rd
AXOR rd=rc,rc	\Leftrightarrow AZERO rd
AADD rd=rc,rc	\Leftrightarrow ASHL.I rd=rc,1
AADD.O rd=rc,rc	\Leftrightarrow ASHL.I.O rd=rc,1
AADD.U.O rd=rc,rc	\Leftrightarrow ASHL.I.U.O rd=rc,1

Selection

class	operation	check
arithmetic	ADD	NONE O U.O
bitwise	OR AND XOR ANDN NOR NAND XNOR ORN	

Format

op rd=rc,rb

rd=op(rc,rb)

31	24 23	18 17	12 11	6 5	0
A.MINOR	rd	rc	rb	op	
8	6	6	6	6	6

Description

The contents of registers rc and rb are fetched and the specified operation is performed on these operands. The result is placed into register rd.

Definition

```

def Address(op,rd,rc,rb) as
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 64)
    case op of
        AADD:
            a ← c + b
        AADD.O:
            t ← {c63 || c} + {b63 || b}
            if t64 ≠ t63 then
                raise FixedPointArithmetic
            endif
            a ← t63..0
        AADD.UO:
            t ← {0¹ || c} + {0¹ || b}
            if t64 ≠ 0 then
                raise FixedPointArithmetic
            endif
            a ← t63..0
        AAND:
            a ← c and b
        AOR:
            a ← c or b
        AXOR:
            a ← c xor b:
        AANDN:
            a ← c and not b
        ANAND:
            a ← not (c and b)
        ANOR:
            a ← not (c or b)
        AXNOR:
            a ← not (c xor b)
        AORN:
            a ← c or not b
    endcase

```

```
RegWrite(rd, 64, a)
enddef
```

Exceptions

Fixed-point arithmetic

Address Compare

These operations perform calculations with two general register values and generate a fixed-point arithmetic exception if the condition specified is met.

Operation codes

A.COM.AND.E	Address compare and equal zero
A.COM.AND.NE	Address compare and not equal zero
A.COM.E	Address compare equal
A.COM.GE	Address compare greater equal signed
A.COM.GE.U	Address compare greater equal unsigned
A.COM.L	Address compare less signed
A.COM.LU	Address compare less unsigned
A.COM.NE	Address compare not equal

Equivalencies

A.COM.E.Z	Address compare equal zero
A.COM.G.Z	Address compare greater zero signed
A.COM.GE.Z	Address compare greater equal zero signed
A.COM.L.Z	Address compare less zero signed
A.COM.LE.Z	Address compare less equal zero signed
A.COM.NE.Z	Address compare not equal zero
A.COM.G	Address compare greater signed
A.COM.G.U	Address compare greater unsigned
A.COM.LE	Address compare less equal signed
A.COM.LE.U	Address compare less equal unsigned
A.FIX	Address fixed point arithmetic exception
A.NOP	Address no operation

A.COM.E.Z rc	\leftarrow A.COM.AND.E rc,rc
A.COM.G.Z rc	\Leftarrow A.COM.L.U rc,rc
A.COM.GE.Z rc	\Leftarrow A.COM.GE rc,rc
A.COM.L.Z rc	\Leftarrow A.COM.L rc,rc
A.COM.LE.Z rc	\Leftarrow A.COM.GE.U rc,rc
A.COM.NE.Z rc	\leftarrow A.COM.AND.NE rc,rc
A.COM.G rc,rd	\rightarrow A.COM.L rd,rc
A.COM.G.U rc,rd	\rightarrow A.COM.L.U rd,rc
A.COM.LE rc,rd	\rightarrow A.COM.GE rd,rc
A.COM.LE.U rc,rd	\rightarrow A.COM.GE.U rd,rc
A.FIX	\leftarrow A.COM.E 0,0
A.NOP	\leftarrow A.COM.NE 0,0

Redundancies


A.COM.E rd,rd	$\Leftrightarrow A.FIX$
A.COM.NE rd,rd	$\Leftrightarrow A.NOP$

Selection

class	operation	cond	operand
boolean	COMAND COM	E NE	
arithmetic	COM	L GE G LE	NONE U
	COM	L GE G LE E NE	Z

Format

A.COM.op rd,rc

acomop(rd,rc)
acomopz(rcd)Description

The contents of registers rd and rc are fetched and the specified condition is calculated on these operands. If the specified condition is true, a fixed-point arithmetic exception is generated. This instruction generates no general register results.

Definition

```

def AddressCompare(op,rd,rc) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    case op of
        A.COM.E:
            a ← d = c
        A.COM.NE:
            a ← d ≠ c
        A.COM.AND.E:
            a ← (d and c) = 0
        A.COM.AND.NE:
            a ← (d and c) ≠ 0
        A.COM.L:
            a ← (rd = rc) ? (c < 0) : (d < c)
        A.COM.GE:
            a ← (rd = rc) ? (c ≥ 0) : (d ≥ c)
        A.COM.LU:
            a ← (rd = rc) ? (c > 0) : ((0 11 d) < (0 11 c))
        A.COM.GE.U:
            a ← (rd = rc) ? (c ≤ 0) : ((0 11 d) ≥ (0 11 c))
    endcase

```

```
if a then
    raise FixedPointArithmetic
endif
enddef
```

Exceptions

Fixed-point arithmetic

Address Copy Immediate

This operation produces one immediate value, placing the result in a general register.

Operation codes

ACOPY.I	Address copy immediate
---------	------------------------

Equivalencies


ASET	Address set
AZERO	Address zero

ASET rd	\leftarrow ACOPY.I rd=-1
AZERO rd	\leftarrow ACOPY.I rd=0

Format

ACOPY.I rd=imm

rd=acopyi(imm)



Description

An immediate value is sign-extended from the 18-bit imm field. The result is placed into register rd.

Definition

```
def AddressCopyImmediate(op,rd,imm) as
    a ← {imm}^10 11 imm
    RegWrite(rd, 128, a)
endef
```

Exceptions

none

Address Immediate

These operations perform calculations with one general register value and one immediate value, placing the result in a general register.

Operation codes

AADD.I	Address add immediate
AADD.I.O	Address add immediate signed check overflow
AADD.I.U.O	Address add immediate unsigned check overflow
AAND.I	Address and immediate
A.NAND.I	Address not and immediate
A.NOR.I	Address not or immediate
A.OR.I	Address or immediate
AXOR.I	Address xor immediate

Equivalencies

AANDN.I	Address and not immediate
ACOPY	Address copy
ANOT	Address not
AORN.I	Address or not immediate
AXNOR.I	Address xnor immediate

AANDN.I rd=rc,imm	→ AAND.I rd=rc,~imm
ACOPY rd=rc	← A.OR.I rd=rc,0
ANOT rd=rc	← ANOR.I rd=rc,0
AORN.I rd=rc,imm	→ A.OR.I rd=rc,~imm
AXNOR.I rd=rc,imm	→ AXOR.I rd=rc,~imm

Redundancies

AADD.I rd=rc,0	↔ ACOPY rd=rc
AADD.I.O rd=rc,0	↔ ACOPY rd=rc
AADD.I.U.O rd=rc,0	↔ ACOPY rd=rc
AAND.I rd=rc,0	↔ AZERO rd
AAND.I rd=rc,-1	↔ ACOPY rd=rc
ANAND.I rd=rc,0	↔ ASET rd
ANAND.I rd=rc,-1	↔ ANOT rd=rc
A.OR.I rd=rc,-1	↔ ASET rd
ANOR.I rd=rc,-1	↔ AZERO rd
AXOR.I rd=rc,0	↔ ACOPY rd=rc
AXOR.I rd=rc,-1	↔ ANOT rd=rc


Selection

class	operation	check
arithmetic	ADD	NONE O UO
bitwise	AND OR NAND NOR XOR	

Format

op rd=rc,imm

rd=op(rc,imm)

Description

The contents of register rc is fetched, and a 64-bit immediate value is sign-extended from the 12-bit imm field. The specified operation is performed on these operands. The result is placed into register rd.

Definition

```
def AddressImmediate(op,rd,rc,imm) as
    i ← imm[11..0]
    c ← RegRead(rc, 64)
    case op of
        AAND.I:
            a ← c and i
        AOR.I:
            a ← c or i
        ANAND.I:
            a ← c nand i
        ANOR.I:
            a ← c nor i
        AXOR.I:
            a ← c xor i
        AADD.I:
            a ← c + i
        AADD.I.O:
            t ← (c63 || c) + (i63 || i)
            if t64 ≠ t63 then
                raise FixedPointArithmetic
            endif
            a ← t63..0
        AADD.I.U.O:
            t ← (c63 || c) + (i63 || i)
            if t64 ≠ 0 then
                raise FixedPointArithmetic
            endif
            a ← t63..0
```

```
endcase
RegWrite(rd, 64, a)
enddef
```

Exceptions

Fixed-point arithmetic

Address Immediate Reversed

These operations perform calculations with one general register value and one immediate value, placing the result in a general register.

Operation codes

A.SET.AND.E.I	Address set and equal immediate
A.SET.AND.NE.I	Address set and not equal immediate
A.SET.E.I	Address set equal immediate
A.SET.GE.I	Address set greater equal immediate signed
A.SET.L.I	Address set less immediate signed
A.SET.NE.I	Address set not equal immediate
A.SET.GE.I.U	Address set greater equal immediate unsigned
A.SET.L.I.U	Address set less immediate unsigned
A.SUB.I	Address subtract immediate
A.SUB.I.O	Address subtract immediate signed check overflow
A.SUB.I.U.O	Address subtract immediate unsigned check overflow

Equivalencies

A.NEG	Address negate
A.NEG.O	Address negate signed check overflow
A.SET.G.I.U	Address set greater immediate unsigned
A.SET.LE.I	Address set less equal immediate signed
A.SET.LE.I.U	Address set less equal immediate unsigned

A.NEG rd=rc	→ A.SUB.I rd=0,rc
A.NEG.O rd=rc	→ A.SUB.I.O rd=0,rc
A.SET.G.I rd=imm,rc	→ A.SET.GE.I rd=imm+1,rc
A.SET.G.I.U rd=imm,rc	→ A.SET.GE.I.U rd=imm+1,rc
A.SET.LE.I rd=imm,rc	→ A.SET.L.I rd=imm-1,rc
A.SET.LE.I.U rd=imm,rc	→ A.SET.L.I.U rd=imm-1,rc

Redundancies

A.SET.AND.E.I rd=rc,0	$\Leftrightarrow A.SET\ rd$
A.SET.AND.NE.I rd=rc,0	$\Leftrightarrow A.ZERO\ rd$
A.SET.AND.E.I rd=rc,-1	$\Leftrightarrow A.SET.E.Z\ rd=rc$
A.SET.AND.NE.I rd=rc,-1	$\Leftrightarrow A.SET.NE.Z\ rd=rc$
A.SET.E.I rd=rc,0	$\Leftrightarrow A.SET.E.Z\ rd=rc$
A.SET.GE.I rd=rc,0	$\Leftrightarrow A.SET.GE.Z\ rd=rc$
A.SET.L.I rd=rc,0	$\Leftrightarrow A.SET.L.Z\ rd=rc$
A.SET.NE.I rd=rc,0	$\Leftrightarrow A.SET.NE.Z\ rd=rc$
A.SET.GE.I.U rd=rc,0	$\Leftrightarrow A.SET.GE.U.Z\ rd=rc$
A.SET.L.I.U rd=rc,0	$\Leftrightarrow A.SET.L.U.Z\ rd=rc$


Selection

class	operation	cond	form	type	check
arithmetic	SUB		I		
				NONE U	O
boolean	SET.AND.SET	E NE	I		
	SET	L GE G LE	I	NONE U	

Format

op rd=imm,rc

rd=op(imm,rc)

Description

The contents of register rc is fetched, and a 64-bit immediate value is sign-extended from the 12-bit imm field. The specified operation is performed on these operands. The result is placed into register rd.

Definition

```
def AddressImmediate(op,rd,rc,imm) as
    i ← imm{11} imm
    c ← RegRead(rc, 64)
    case op of
        ASUB.I:
            a ← i - c
        ASUB.I.O:
            t ← (i63 || i) - (c63 || c)
            if t64 ≠ t3 then
                raise FixedPointArithmetic
            endif
```

```

    a ← t63..0
ASUB.I.U.O:
    t ← (i63 .. i) - (c63 .. c)
    if t64 ≠ 0 then
        raise FixedPointArithmetic
    endif
    a ← t63..0
ASET.AND.E.I:
    a ← (i and c) = 064
ASET.AND.NE.I:
    a ← (i and c) ≠ 064
ASET.E.I:
    a ← (i = c)64
ASET.NE.I:
    a ← (i ≠ c)64
ASET.L.I:
    a ← (i < c)64
ASET.GE.I:
    a ← (i ≥ c)64
ASET.L.I.U:
    a ← ((0 .. i) < (0 .. c))64
ASET.GE.I.U:
    a ← ((0 .. i) ≥ (0 .. c))64
endcase
RegWrite(rd, 64, a)
enddef

```

Exceptions

Fixed-point arithmetic

Address Reversed

These operations perform calculations with two general register values, placing the result in a general register.

Operation codes

A.SET.AND.E	Address set and equal zero
A.SET.AND.NE	Address set and not equal zero
A.SET.E	Address set equal
A.SET.GE	Address set greater equal signed
A.SET.GE.U	Address set greater equal unsigned
A.SET.L	Address set less signed
A.SET.L.U	Address set less unsigned
A.SET.NE	Address set not equal
A.SUB	Address subtract
A.SUB.O	Address subtract signed check overflow
A.SUB.U.O	Address subtract unsigned check overflow

Equivalencies

A.SET.E.Z	Address set equal zero
A.SET.G.Z	Address set greater zero signed
A.SET.GE.Z	Address set greater equal zero signed
A.SET.L.Z	Address set less zero signed
A.SET.LE.Z	Address set less equal zero signed
A.SET.NE.Z	Address set not equal zero
A.SET.G	Address set greater signed
A.SET.G.U	Address set greater unsigned
A.SET.LE	Address set less equal signed
A.SET.LE.U	Address set less equal unsigned

A.SET.E.Z rd=rc	← A.SET.AND.E rd=rc,rc
A.SET.G.Z rd=rc	← A.SET.L.U rd=rc,rc
A.SET.GE.Z rd=rc	← A.SET.GE rd=rc,rc
A.SET.L.Z rd=rc	← A.SET.L rd=rc,rc
A.SET.LE.Z rd=rc	← A.SET.GE.U rd=rc,rc
A.SET.NE.Z rd=rc	← A.SET.AND.NE rd=rc,rc
A.SET.G rd=rb,rc	→ A.SET.L rd=rc,rb
A.SET.G.U rd=rb,rc	→ A.SET.L.U rd=rc,rb
A.SET.LE rd=rb,rc	→ A.SET.GE rd=rc,rb
A.SET.LE.U rd=rb,rc	→ A.SET.GE.U rd=rc,rb

Redundancies


A.SET.E rd=rc,rc	\Leftrightarrow	A.SET rd
A.SET.NE rd=rc,rc	\Leftrightarrow	A.ZERO rd

Selection

class	operation	cond	operand	check
arithmetic	SUB			
boolean	SET.AND SET	E NE		
	SET	L GE G LE	NONE U	
	SET	L GE G LE E NE	Z	

Format

op rd=rb,rc

rd=op(rb,rc)
rd=opz(rcb)rc \leftarrow rb \leftarrow rcbDescription

The contents of registers rc and rb are fetched and the specified operation is performed on these operands. The result is placed into register rd.

Definition

```
def AddressReversed(op,rd,rc,rb) as
    c <- RegRead[rc, 128]
    b <- RegRead[rb, 128]
    case op of
        A.SET.E:
            a <- (b = c)^64
        A.SET.NE:
            a <- (b ≠ c)^64
        A.SET.AND.E:
            a <- ((b and c) = 0)^64
        A.SET.AND.NE:
            a <- ((b and c) ≠ 0)^64
        A.SET.L:
            a <- ((rc = rb) ? (b < 0) : (b < c))^64
        A.SET.GE:
            a <- ((rc = rb) ? (b ≥ 0) : (b ≥ c))^64
        A.SET.L.U:
```

```

 $a \leftarrow ((rc = rb) ? (b > 0) : (|0 \mid\mid b| < |0 \mid\mid c|))^{64}$ 
ASET.GE.U:
 $a \leftarrow ((rc = rb) ? (b \leq 0) : (|0 \mid\mid b| \geq |0 \mid\mid c|))^{64}$ 
ASUB:
 $a \leftarrow b - c$ 
ASUB.O:
 $t \leftarrow (b_{63} \mid\mid b) - (c_{63} \mid\mid c)$ 
if  $t_{64} \neq t_{63}$  then
    raise FixedPointArithmetic
endif
 $a \leftarrow t_{63..0}$ 
ASUB.U.O:
 $t \leftarrow (0^1 \mid\mid b) - (0^1 \mid\mid c)$ 
if  $t_{64} \neq 0$  then
    raise FixedPointArithmetic
endif
 $a \leftarrow t_{63..0}$ 
endcase
RegWrite(rd, 64, a)
enddef

```

Exceptions

Fixed-point arithmetic

Address Shift Left Immediate Add

These operations perform calculations with two general register values, placing the result in a general register.


Operation codes

ASHLI.ADD	Address shift left immediate add
-----------	----------------------------------

Format

ASHLI.ADD rd=rc,rb,i

rc=op(ra,rb,i)



assert $1 \leq i \leq 4$

sh $\leftarrow i-1$

Description

The contents of register rb are shifted left by the immediate amount and added to the contents of register rc. The result is placed into register rd.

Definition

```
def AddressShiftLeftImmediateAdd(sh,rd,rc,rb) as
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 64)
    a ← c + (b62-sh..0 || 01+sh)
    RegWrite(rd, 64, a)
endef
```

Exceptions

none

Address Shift Left Immediate Subtract

These operations perform calculations with two general register values, placing the result in a general register.

Operation codes

ASHLI.SUB	Address shift left immediate subtract
-----------	---------------------------------------

Format

ASHLI.SUB rd=rb,i,rc

rd=op(rb,i,rc)

31	24 23	18 17	12 11	6 5	21 0
A.MINOR	rd	rc	rb	ASHLI.SUB	sh

assert $1 \leq i \leq 4$

sh $\leftarrow i-1$

Description

The contents of register rc is subtracted from the contents of register rb shifted left by the immediate amount. The result is placed into register rd.

Definition

```
def AddressShiftLeftImmediateSubtract{op,rd,rc,rb} as
    c ← RegRead[rc, 128]
    b ← RegRead[rb, 128]
    a ← (b62-sh..0 11 01+sh) - c
    RegWrite[rd, 64, a]
enddef
```

Exceptions

none

Address Shift Immediate

These operations perform calculations with one general register value and one immediate value, placing the result in a general register.

Operation codes

A.SHL.I	Address shift left immediate
A.SHL.I.O	Address shift left immediate signed check overflow
A.SHL.I.U.O	Address shift left immediate unsigned check overflow
A.SHR.I	Address signed shift right immediate
A.SHR.I.U	Address shift right immediate unsigned

Redundancies

A.SHL.I rd=rc,1	\Leftrightarrow	AADD rd=rc,rc
A.SHL.I.O rd=rc,1	\Leftrightarrow	AADD.O rd=rc,rc
A.SHL.I.U.O rd=rc,1	\Leftrightarrow	AADD.U.O rd=rc,rc
A.SHL.I rd=rc,0	\Leftrightarrow	ACOPY rd=rc
A.SHL.I.O rd=rc,0	\Leftrightarrow	ACOPY rd=rc
A.SHL.I.U.O rd=rc,0	\Leftrightarrow	ACOPY rd=rc
A.SHR.I rd=rc,0	\Leftrightarrow	ACOPY rd=rc
A.SHR.I.U rd=rc,0	\Leftrightarrow	ACOPY rd=rc


Selection

class	operation	form	operand	check
shift	SHL	I		
			NONE U	O
	SHR	I	NONE U	

Format

op rd=rc,simm

rd=op(rc,simm)



Description

The contents of register rc is fetched, and a 6-bit immediate value is taken from the 6-bit simm field. The specified operation is performed on these operands. The result is placed into register rd.

Definition

```

def AddressShiftImmediate(op,rd,rc,simm) as
    c ← RegRead(rc, 64)
    case op of
        ASHL.I:
            a ← c63.simm..0 || 0simm
        ASHL.I.O:
            if c63..63-simm ≠ c63simm+1 then
                raise FixedPointArithmetic
            endif
            a ← c63.simm..0 || 0simm
        ASHL.I.U:
            if c63..64-simm ≠ 0 then
                raise FixedPointArithmetic
            endif
            a ← c63.simm..0 || 0simm
        ASHR.I:
            a ← a63simm || c63.simm
        ASHR.I.U:
            a ← 0simm || c63.simm
    endcase
    RegWrite(rd, 64, a)
enddef

```

Exceptions

Fixed-point arithmetic

Address Ternary

These operations perform calculations with three general register values, placing the result in a fourth general register.


Operation codes

A.MUX	Address multiplex
-------	-------------------

Format

op ra=rd,rc,rb

ra=amux(rd,rc,rb)



Description

The contents of registers rd, rc, and rb are fetched. The specified operation is performed on these operands. The result is placed into register ra.

Definition

```
def AddressTernary(op,rd,rc,rb,ra) as
    d ← RegRead(rd, 64)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 64)
    endcase
    case op of
        A.MUX:
            a ← (c and d) or (b and not d)
        endcase
        RegWrite(ra, 64, a)
    enddef
```

Exceptions

none

Branch


This operation branches to a location specified by a register.

Operation codes

B	Branch
---	--------

Format

B rd



Description

Execution branches to the address specified by the contents of register rd.

Access disallowed exception occurs if the contents of register rd is not aligned on a quadlet boundary.

Definition

```
def Branch(rd,rc,rb) as
    if {rc ≠ 0} or {rb ≠ 0} then
        raise ReservedInstruction
    endif
    d ← RegRead{rd, 64}
    if {d1..0} ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    endif
    ProgramCounter ← d63..2 || 02
    raise TakenBranch
enddef
```

Exceptions

Reserved Instruction
Access disallowed by virtual address

Branch Back

This operation branches to a location specified by the previous contents of register 0, reduces the current privilege level, loads a value from memory, and restores register 0 to the value saved on a previous exception.

Operation codes

B.BACK	Branch back
--------	-------------

Format

B.BACK

bback()

31	24 23	18 17	12 11	6 5	0
	B.MINOR	0	0	0	B.BACK

8 6 6 6 6

Description

Processor context, including program counter and privilege level is restored from register 0, where it was saved at the last exception. Exception state, if set, is cleared, re-enabling normal exception handling. The contents of register 0 saved at the last exception is restored from memory. The privilege level is only lowered, so that this instruction need not be privileged.

If the previous exception was an AccessDetail exception, Continuation State set at the time of the exception affects the operation of the next instruction after this Branch Back, causing the previous AccessDetail exception to be inhibited. If software is performing this instruction to abort a sequence ending in an AccessDetail exception, it should abort by branching to an instruction that is not affected by Continuation State.

Definition

```

def BranchBack(rd,rc,rb) as
    c ← RegRead[rc, 128]
    if (rd ≠ 0) or (rc ≠ 0) or (rb ≠ 0) then
        raise ReservedInstruction
    endif
    a ← LoadMemory(ExceptionBase,ExceptionBase+Thread*128,128,L)
    if PrivilegeLevel > c1..0 then
        PrivilegeLevel ← c1..0
    endif
    ProgramCounter ← c63..2 || 02
    ExceptionState ← 0
    RegWrite[rd,128,a]
    raise TakenBranchContinue
enddef

```

Exceptions

Reserved Instruction

Access disallowed by virtual address

Access disallowed by tag

Access disallowed by global TB

Access disallowed by local TB

Access detail required by tag

Access detail required by local TB

Access detail required by global TB

Local TB miss

Global TB miss

Branch Barrier

This operation stops the current thread until all pending stores are completed, then branches to a location specified by a register.


Operation codes

B.BARRIER	Branch barrier
-----------	----------------

Format

B.BARRIER rd

bbarrier(rd)



Description

The instruction fetch unit is directed to cease execution until all pending stores are completed. Following the barrier, any previously pre-fetched instructions are discarded and execution branches to the address specified by the contents of register rd.

Access disallowed exception occurs if the contents of register rd is not aligned on a quadlet boundary.

Self-modifying, dynamically-generated, or loaded code may require use of this instruction between storing the code into memory and executing the code.

Definition

```
def BranchBarrier(rd,rc,rb) as
    if (rc ≠ 0) or (rb ≠ 0) then
        raise ReservedInstruction
    endif
    d ← RegRead(rd, 64)
    if (d1..0) ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    endif
    ProgramCounter ← d63..2 || 02
    FetchBarrier()
    raise TakenBranch
enddef
```

Exceptions

Reserved Instruction

Branch Conditional

These operations compare two operands, and depending on the result of that comparison, conditionally branches to a nearby code location.

Operation codes

B.AND.E	Branch and equal zero
B.AND.NE	Branch and not equal zero
B.E	Branch equal
B.GE	Branch greater equal signed
B.L	Branch signed less
B.NE	Branch not equal
B.GE.U	Branch greater equal unsigned
B.L.U	Branch less unsigned

Equivalencies

B.E.Z	Branch equal zero
B.G.Z ⁶	Branch greater zero signed
B.GE.Z ⁷	Branch greater equal zero signed
B.L.Z ⁸	Branch less zero signed
B.LE.Z ⁹	Branch less equal zero signed
B.NE.Z	Branch not equal zero
B.LE	Branch less equal signed
B.G	Branch greater signed
B.LE.U	Branch less equal unsigned
B.G.U	Branch greater unsigned
B.NOP	Branch no operation

⁶B.G.Z is encoded as B.L.U with both instruction fields rd and rc equal.

⁷B.GE.Z is encoded as B.GE with both instruction fields rd and rc equal.

⁸B.L.Z is encoded as B.L with both instruction fields rd and rc equal.

⁹B.LE.Z is encoded as B.GE.U with both instruction fields rd and rc equal.

<i>B.E.Z rc,target</i>	\leftarrow B.AND.E rc,rc,target
<i>B.G.Z rc,target</i>	\Leftarrow B.L.U rc,rc,target
<i>B.GE.Z rc,target</i>	\Leftarrow B.GE rc,rc,target
<i>B.L.Z rc,target</i>	\Leftarrow B.L rc,rc,target
<i>B.LE.Z rc,target</i>	\Leftarrow B.GE.U rc,rc,target
<i>B.NE.Z rc,target</i>	\leftarrow B.AND.NE rc,rc,target
<i>B.LE rc,rd,target</i>	\rightarrow B.GE rd,rc,target
<i>B.G rc,rd,target</i>	\rightarrow B.L rd,rc,target
<i>B.LE.U rc,rd,target</i>	\rightarrow B.GE.U rd,rc,target
<i>B.G.U rc,rd,target</i>	\rightarrow B.L.U rd,rc,target
<i>B.NOP</i>	\leftarrow B.NE r0,r0,\$

Redundancies

<i>B.E rc,rc,target</i>	\Leftrightarrow B.I target
<i>B.NE rc,rc,target</i>	\Leftrightarrow B.NOP


Selection

class	op	compare	type
arithmetic		L GE G LE	NONEU
vs. zero		L GE G LE E NE	Z
bitwise	none AND	E NE	

Format

op rd,rc,target

if {op(rd,rc)} goto target;

Description

The contents of registers rd and rc are compared, as specified by the op field. If the result of the comparison is true, execution branches to the address specified by the offset field. Otherwise, execution continues at the next sequential instruction.

Definition

```
def BranchConditionally(op,rd,rc,offset) as
    d ← RegRead[rd, 128]
    c ← RegRead[rc, 128]
    case op of
        B.E:
```

```

    a ← d = c
B.NE:
    a ← d ≠ c
B.AND.E:
    a ← (d and c) = 0
B.AND.NE:
    a ← (d and c) ≠ 0
B.L:
    a ← (rd = rc) ? (c < 0): (d < c)
B.GE:
    a ← (rd = rc) ? (c ≥ 0): (d ≥ c)
B.LU:
    a ← (rd = rc) ? (c > 0): (|0 || d) < (|0 || c))
B.GE.U:
    a ← (rd = rc) ? (c ≤ 0): (|0 || d) ≥ (|0 || c))
endcase
if a then
    ProgramCounter ← ProgramCounter + (offset5011 offset11 02)
    raise TakenBranch
endif
enddef

```

Exceptions

none

Branch Conditional Floating-Point

These operations compare two floating-point operands, and depending on the result of that comparison, conditionally branches to a nearby code location.

Operation codes

B.E.F.16	Branch equal floating-point half
B.E.F.32	Branch equal floating-point single
B.E.F.64	Branch equal floating-point double
B.E.F.128	Branch equal floating-point quad
B.GE.F.16	Branch greater equal floating-point half
B.GE.F.32	Branch greater equal floating-point single
B.GE.F.64	Branch greater equal floating-point double
B.GE.F.128	Branch greater equal floating-point quad
B.L.F.16	Branch less floating-point half
B.L.F.32	Branch less floating-point single
B.L.F.64	Branch less floating-point double
B.L.F.128	Branch less floating-point quad
B.LG.F.16	Branch less greater floating-point half
B.LG.F.32	Branch less greater floating-point single
B.LG.F.64	Branch less greater floating-point double
B.LG.F.128	Branch less greater floating-point quad

Equivalencies

B.LE.F.16	Branch less equal floating-point half
B.LE.F.32	Branch less equal floating-point single
B.LE.F.64	Branch less equal floating-point double
B.LE.F.128	Branch less equal floating-point quad
B.G.F.16	Branch greater floating-point half
B.G.F.32	Branch greater floating-point single
B.G.F.64	Branch greater floating-point double
B.G.F.128	Branch greater floating-point quad

B.LE.F.size rc,rd,target	→ B.GE.F.size rd,rc,target
B.G.F.size rc,rd,target	→ B.L.F.size rd,rc,target


Selection

number format	type	compare	size
floating-point	F	E LG L GE G LE	16 32 64 128

Format

```
op rd,rc,target
```

```
if (op(rd,rc)) goto target;
```

Description

The contents of registers rc and rd are compared, as specified by the op field. If the result of the comparison is true, execution branches to the address specified by the offset field. Otherwise, execution continues at the next sequential instruction.

Definition

```
def BranchConditional(FloatingPointOp,rd,rc,offset) as
  case op of
    B.E.F.16, B.LG.F.16, B.L.F.16, B.GE.F.16:
      size ← 16
    B.E.F.32, B.LG.F.32, B.L.F.32, B.GE.F.32:
      size ← 32
    B.E.F.64, B.LG.F.64, B.L.F.64, B.GE.F.64:
      size ← 64
    B.E.F.128, B.LG.F.128, B.L.F.128, B.GE.F.128:
      size ← 128
  endcase
  d ← F(size,RegRead(rd, 128))
  c ← F(size,RegRead(rc, 128))
  v ← fcom(d, c)
  case op of
    BEF16, BEF32, BEF64, BEF128:
      a ← {v = E}
    BLGF16, BLGF32, BLGF64, BLGF128:
      a ← {v = L} or {v = G}
    BLF16, BLF32, BLF64, BLF128:
      a ← {v = L}
    BGEF16, BGEF32, BGEF64, BGEF128:
      a ← {v = G} or {v = E}
  endcase
  if a then
    ProgramCounter ← ProgramCounter + (offset50 || offset11 || 02)
    raise TakenBranch
  endif
enddef
```

Exceptions

```
none
```

Branch Conditional Visibility Floating-Point

These operations compare two group-floating-point operands, and depending on the result of that comparison, conditionally branches to a nearby code location.

Operation codes

B.I.F.32	Branch invisible floating-point single
B.NI.F.32	Branch not invisible floating-point single
B.NV.F.32	Branch not visible floating-point single
B.V.F.32	Branch visible floating-point single


Selection

number format	type	compare	size
floating-point	F	I NI NV V	32

Format

op rc,rd,target

if (op(rc,rd)) goto target;



Description

The contents of registers rc and rd are compared, as specified by the op field. If the result of the comparison is true, execution branches to the address specified by the offset field. Otherwise, execution continues at the next sequential instruction.

Each operand is assumed to represent a vertex of the form: [w z y x] packed into a single register. The comparisons check for visibility of a line connecting the vertices against a standard viewing volume, defined by the planes: x=w,x=-w,y=w,y=-w,z=0,z=1. A line is visible (V) if the vertices are both within the volume. A line is not visible (NV) is either vertex is outside the volume - in such a case, the line may be partially visible. A line is invisible (I) if the vertices are both outside any face of the volume. A line is not invisible (NI) if the vertices are not both outside any face of the volume.

Definition

```
def n(a) as {a.t=QNAN} or {a.t=SNAN} enddef
```

```
def less(a,b) as fcom(a,b)=L enddef
```

```
def trxya,b,c,d) as {fcom(fabs(a),b)=G} and {fcom(fabs(c),d)=G} and {a.s=c.s} enddef
```

```

def BranchConditionalVisibilityFloatingPoint(op,rd,rc,offset) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    dx ← F{32,d31..0}
    cx ← F{32,c31..0}
    dy ← F{32,d63..32}
    cy ← F{32,c63..32}
    dz ← F{32,d95..64}
    cz ← F{32,c95..64}
    dw ← F{32,d127..96}
    cw ← F{32,c127..96}
    f1 ← F{32,0x7f000000} // floating-point 1.0
    if (n(dx) or n(dy) or n(dz) or n(dw) or n(cx) or n(cy) or n(cz) or n(cw)) then
        a ← false
    else
        dv ← less(fabs(dx),dz) and less(fabs(dy),dz) and less(dz,f1) and (dz.s=0)
        cv ← less(fabs(cx),cz) and less(fabs(cy),cz) and less(cz,f1) and (cz.s=0)
        trz ← (less(f1,dz) and less(f1,cz)) or ((dz.s=1 and cz.s=1))
        tr ← trxy(dx,dz,cx,cz) or trxy(dy,dz,cy,cz) or trz
        case op of
            B.I.F.32:
                a ← tr
            B.NI.F.32:
                a ← not tr
            B.NV.F.32:
                a ← not (dv and cv)
            B.V.F.32:
                a ← dv and cv
        endcase
    endif
    if a then
        ProgramCounter ← ProgramCounter + (offset50 || offset || 02)
        raise TakenBranch
    endif
enddef

```

Exceptions

none

Branch Down

This operation branches to a location specified by a register, reducing the current privilege level.


Operation codes

B.DOWN	Branch down
--------	-------------

Format

B.DOWN rd

bdown(rd)



Description

Execution branches to the address specified by the contents of register rd. The current privilege level is reduced to the level specified by the low order two bits of the contents of register rd.

Definition

```
def BranchDown(rd,rc,rb) as
    if (rc ≠ 0) or (rb ≠ 0) then
        raise ReservedInstruction
    endif
    d ← RegRead(rd, 64)
    if PrivilegeLevel > d1..0 then
        PrivilegeLevel ← d1..0
    endif
    ProgramCounter ← d63..2 || 02
    raise TakenBranch
enddef
```

Exceptions

Reserved Instruction

Branch Gateway

This operation provides a secure means to call a procedure, including those at a higher privilege level.

Operation codes

B.GATE	Branch gateway
--------	----------------

Equivalencies

B.GATE	← B.GATE 0
--------	------------

Format

B.GATE rb

bgate(rb)

31	24 23	18 17	12 11	6 5	0
B.MINOR	0	1	rb	B.GATE	

8 6 6 6 6

Description

The contents of register rb is a branch address in the high-order 62 bits and a new privilege level in the low-order 2 bits. A branch and link occurs to the branch address, and the privilege level is raised to the new privilege level. The high-order 62 bits of the successor to the current program counter is catenated with the 2-bit current execution privilege and placed in register 0.

If the new privilege level is greater than the current privilege level, an octlet of memory data is fetched from the address specified by register 1, using the little-endian byte order and a gateway access type. A *GatewayDisallowed* exception occurs if the original contents of register 0 do not equal the memory data.


If the new privilege level is the same as the current privilege level, no checking of register 1 is performed.

An *AccessDisallowed* exception occurs if the new privilege level is greater than the privilege level required to write the memory data, or if the old privilege level is lower than the privilege required to access the memory data as a gateway, or if the access is not aligned on an 8-byte boundary.

A *ReservedInstruction* exception occurs if the rc field is not one or the rd field is not zero.

In the example below, a gateway from level 0 to level 2 is illustrated. The gateway pointer, located by the contents of register rc (1), is fetched from memory and compared against the

contents of register rb (0). The instruction may only complete if these values are equal. Concurrently, the contents of register rb (0) is placed in the program counter and privilege level, and the address of the next sequential address and privilege level is placed into register rd (0). Code at the target of the gateway locates the data pointer at an offset from the gateway pointer (register 1), and fetches it into register 1, making a data region available. A stack pointer may be saved and fetched using the data region, another region located from the data region, or a data region located as an offset from the original gateway pointer.



Branch gateway

For additional information on the branch-gateway instruction, see the [System and Privileged Library Calls](#) section on page 44.

This instruction gives the target procedure the assurances that register 0 contains a valid return address and privilege level, that register 1 points to the gateway location, and that the gateway location is octet aligned. Register 1 can then be used to

securely reach values in memory. If no sharing of literal pools is desired, register 1 may be used as a literal pool pointer directly. If sharing of literal pools is desired, register 1 may be used with an appropriate offset to load a new literal pool pointer; for example, with a one cache line offset from the register 1. Note that because the virtual memory system operates with cache line granularity, that several gateway locations must be created together.

Software must ensure that an attempt to use any octlet within the region designated by virtual memory as gateway either functions properly or causes a legitimate exception. For example, if the adjacent octlets contain pointers to literal pool locations, software should ensure that these literal pools are not executable, or that by virtue of being aligned addresses, cannot raise the execution privilege level. If register 1 is used directly as a literal pool location, software must ensure that the literal pool locations that are accessible as a gateway do not lead to a security violation.

Register 0 contains a valid return address and privilege level, the value is suitable for use directly in the Branch-down (B.DOWN) instruction to return to the gateway callee.

Definition

```
def BranchGateway(rd,rc,rb) as
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 64)
    if (rd ≠ 0) or (rc ≠ 1) then
        raise ReservedInstruction
    endif
    if c2..0 ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    endif
    d ← ProgramCounter63..2+1 || PrivilegeLevel
    if PrivilegeLevel < b1..0 then
        m ← LoadMemoryG(c,c,64,L)
        if b ≠ m then
            raise GatewayDisallowed
        endif
        PrivilegeLevel ← b1..0
    endif
    ProgramCounter ← b63..2 || 02
    RegWrite(rd, 64, d)
    raise TakenBranch
enddef
```

Exceptions

- Reserved Instruction
- Gateway disallowed
- Access disallowed by virtual address
- Access disallowed by tag
- Access disallowed by global TB
- Access disallowed by local TB
- Access detail required by tag

Access detail required by local TB
Access detail required by global TB
Local TB miss
Global TB miss

Branch Halt

This operation stops the current thread until an exception occurs.

Operation codes

B.HALT	Branch halt
--------	-------------

Format

B.HALT

bhalt()

31	24 23	18 17	12 11	6 5	0
	B.MINOR	0	0	0	B.HALT

8 6 6 6 6

Description

This instruction directs the instruction fetch unit to cease execution until an exception occurs.

Definition

```
def BranchHalt(rd,rc,rb) as
    if (rd ≠ 0) or (rc ≠ 0) or (rb ≠ 0) then
        raise ReservedInstruction
    endif
    FetchHalt()
enddef
```

Exceptions

Reserved Instruction

Branch Hint

This operation indicates a future branch location specified by a register.

Operation codes

B.HINT	Branch Hint
--------	-------------

Format

B.HINT badd,count,rd

bhint(badd,count,rd)

31	24 23	18 17	12 11	6 5	0
	B.MINOR	rd	count	simm	B.HINT

8 6 6 6 6

simm \leftarrow badd-pc-4

Description

This instruction directs the instruction fetch unit of the processor that a branch is likely to occur count times at simm instructions following the current successor instruction to the address specified by the contents of register rd.

After branching count times, the instruction fetch unit should presume that the branch at simm instructions following the current successor instruction is not likely to occur. If count is zero, this hint directs the instruction fetch unit that the branch is likely to occur more than 63 times.

Access disallowed exception occurs if the contents of register rd is not aligned on a quadlet boundary.

Definition

```
def BranchHint(rd,count,simm) as
    d  $\leftarrow$  RegRead(rd, 64)
    if  $|d_{1..0}| \neq 0$  then
        raise AccessDisallowedByVirtualAddress
    endif
    FetchHint(ProgramCounter +4 + (0 11 simm 11 02), d63..2 11 02, count)
endef
```

Exceptions

Access disallowed by virtual address

Branch Hint Immediate

This operation indicates a future branch location specified as an offset from the program counter.


Operation codes

B.HINT.I	Branch Hint Immediate
----------	-----------------------

Format

B.HINT.I badd,count,target

bhinti(badd,count,target)



simm \leftarrow badd-pc-4

Description

This instruction directs the instruction fetch unit of the processor that a branch is likely to occur **count** times at **simm** instructions following the current successor instruction to the address specified by the **offset** field.

After branching **count** times, the instruction fetch unit should presume that the branch at **simm** instructions following the current successor instruction is not likely to occur. If **count** is zero, this hint directs the instruction fetch unit that the branch is likely to occur more than 63 times.

Definition

```
def BranchHintImmediate(simm,count,offset) as
    BranchHint(ProgramCounter + 4 + {0 11 simm 11 02}, count,
               ProgramCounter + {offset 11 11 offset 11 02})
enddef
```

Exceptions

none

Branch Immediate

This operation branches to a location that is specified as an offset from the program counter.

Operation codes

B.I	Branch immediate
-----	------------------


Redundancies

B.I target	\Leftrightarrow	B.E rc,rc,target
------------	-------------------	------------------

Format

B.I target

bi[target]



Description

Execution branches to the address specified by the offset field.

Definition

```
def BranchImmediate(offset) as
    ProgramCounter ← ProgramCounter + (offset23 || offset11 || 02)
    raise TakenBranch
enddef
```

Exceptions

none

Branch Immediate Link

This operation branches to a location that is specified as an offset from the program counter, saving the value of the program counter into register 0.


Operation codes

B.LINK.I	Branch immediate link
----------	-----------------------

Format

B.LINK.I target

blinki(target)



Description

The address of the instruction following this one is placed into register 0. Execution branches to the address specified by the offset field.

Definition

```
def BranchImmediateLink(offset) as
    RegWrite(0, 64, ProgramCounter +. 4)
    ProgramCounter ← ProgramCounter + (offset2338 || offset || 02)
    raise TakenBranch
enddef
```

Exceptions

none

Branch Link

This operation branches to a location specified by a register, saving the value of the program counter into a register.

Operation codes


B.LINK	Branch link
--------	-------------

Equivalencies

B.LINK	\leftarrow B.LINK 0=0
B.LINK rc	\leftarrow B.LINK 0=rc

Format

B.LINK rd=rc



rb \leftarrow 0

Description

The address of the instruction following this one is placed into register rd. Execution branches to the address specified by the contents of register rc.

Access disallowed exception occurs if the contents of register rc is not aligned on a quadlet boundary.

Reserved instruction exception occurs if rb is not zero.

Definition

```
def BranchLink(rd,rc,rb) as
    if rb  $\neq$  0 then
        raise ReservedInstruction
    endif
    c  $\leftarrow$  RegRead(rc, 64)
    if (c and 3)  $\neq$  0 then
        raise AccessDisallowedByVirtualAddress
    endif
    RegWrite(rd, 64, ProgramCounter + 4)
    ProgramCounter  $\leftarrow$  c63..2 || 02
    raise TakenBranch
enddef
```

Exceptions

Reserved Instruction

Access disallowed by virtual address

Load

These operations compute a virtual address from the contents of two registers, load data from memory, sign- or zero-extending the data to fill the destination register.

Operation codes

L.8 ¹⁰	Load signed byte
L.16.B	Load signed doublet big-endian
L.16.AB	Load signed doublet aligned big-endian
L.16.L	Load signed doublet little-endian
L.16.AL	Load signed doublet aligned little-endian
L.32.B	Load signed quadlet big-endian
L.32.AB	Load signed quadlet aligned big-endian
L.32.L	Load signed quadlet little-endian
L.32.AL	Load signed quadlet aligned little-endian
L.64.B	Load signed octlet big-endian
L.64.AB	Load signed octlet aligned big-endian
L.64.L	Load signed octlet little-endian
L.64.AL	Load signed octlet aligned little-endian
L.128.B ¹¹	Load hexlet big-endian
L.128.AB ¹²	Load hexlet aligned big-endian
L.128.L ¹³	Load hexlet little-endian
L.128.AL ¹⁴	Load hexlet aligned little-endian
L.U.8 ¹⁵	Load unsigned byte
L.U.16.B	Load unsigned doublet big-endian
L.U.16.AB	Load unsigned doublet aligned big-endian
L.U.16.L	Load unsigned doublet little-endian
L.U.16.AL	Load unsigned doublet aligned little-endian
L.U.32.B	Load unsigned quadlet big-endian
L.U.32.AB	Load unsigned quadlet aligned big-endian
L.U.32.L	Load unsigned quadlet little-endian
L.U.32.AL	Load unsigned quadlet aligned little-endian
L.U.64.B	Load unsigned octlet big-endian
L.U.64.AB	Load unsigned octlet aligned big-endian
L.U.64.L	Load unsigned octlet little-endian
L.U.64.AL	Load unsigned octlet aligned little-endian

¹⁰L.8 need not distinguish between little-endian and big-endian ordering, nor between aligned and unaligned, as only a single byte is loaded.

¹¹L.128.B need not distinguish between signed and unsigned, as the hexlet fills the destination register.

¹²L.128.AB need not distinguish between signed and unsigned, as the hexlet fills the destination register.

¹³L.128.L need not distinguish between signed and unsigned, as the hexlet fills the destination register.

¹⁴L.128.AL need not distinguish between signed and unsigned, as the hexlet fills the destination register.

¹⁵L.U8 need not distinguish between little-endian and big-endian ordering, nor between aligned and unaligned, as only a single byte is loaded.


Selection

number format	type	size	alignment	ordering
signed byte		8		
unsigned byte	U	8		
signed integer		16 32 64		L B
signed integer aligned		16 32 64	A	L B
unsigned integer	U	16 32 64		L B
unsigned integer aligned	U	16 32 64	A	L B
register		128		L B
register aligned		128	A	L B

Format

op rd=rc,rb

rd=op(rc,rb)

Description

An operand size, expressed in bytes, is specified by the instruction. A virtual address is computed from the sum of the contents of register rc and the contents of register rb multiplied by operand size. The contents of memory using the specified byte order are read, treated as the size specified, zero-extended or sign-extended as specified, and placed into register rd.

If alignment is specified, the computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an "access disallowed by virtual address" exception occurs.

Definition

```

def Load{op,rd,rc,rb} as
  case op of
    L16L, L32L, L8, L16AL, L32AL, L16B, L32B, L16AB, L32AB,
    L64L, L64AL, L64B, L64AB:
      signed ← true
    LU16L, LU32L, LU8, LU16AL, LU32AL, LU16B, LU32B, LU16AB, LU32AB,
    LU64L, LU64AL, LU64B, LU64AB:
      signed ← false
    L128L, L128AL, L128B, L128AB:
      signed ← undefined
  endcase
  case op of
    L8, LU8:
      size ← 8
    L16L, LU16L, L16AL, LU16AL, L16B, LU16B, L16AB, LU16AB:
  endcase

```

```

        size ← 16
L32L, LU32L, L32AL, LU32AL, L32B, LU32B, L32AB, LU32AB:
        size ← 32
L64L, LU64L, L64AL, LU64AL, L64B, LU64B, L64AB, LU64AB:
        size ← 64
L128L, L128AL, L128B, L128AB:
        size ← 128
endcase
lsize ← log(size)
case op of
    L16L, LU16L, L32L, LU32L, L64L, LU64L, L128L,
    L16AL, LU16AL, L32AL, LU32AL, L64AL, LU64AL, L128AL:
        order ← L
    L16B, LU16B, L32B, LU32B, L64B, LU64B, L128B,
    L16AB, LU16AB, L32AB, LU32AB, L64AB, LU64AB, L128AB:
        order ← B
    L8, LU8:
        order ← undefined
endcase
c ← RegRead(rc, 64)
b ← RegRead(rb, 64)
VirtAddr ← c + (b66-lsize..0 || 0lsize-3)
case op of
    L16AL, LU16AL, L32AL, LU32AL, L64AL, LU64AL, L128AL,
    L16AB, LU16AB, L32AB, LU32AB, L64AB, LU64AB, L128AB:
        if {csize-4..0 ≠ 0 then
            raise AccessDisallowedByVirtualAddress
        endif
    L16L, LU16L, L32L, LU32L, L64L, LU64L, L128L,
    L16B, LU16B, L32B, LU32B, L64B, LU64B, L128B:
    L8, LU8:
endcase
m ← LoadMemory(c, VirtAddr, size, order)
a ← (msize-1 and signed)128-size || m
RegWrite(rd, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Load Immediate

These operations compute a virtual address from the contents of a register and a sign-extended immediate value, load data from memory, sign- or zero-extending the data to fill the destination register.

Operation codes

L.I.8 ¹⁶	Load immediate signed byte
L.I.16.A.B	Load immediate signed doublet aligned big-endian
L.I.16.B	Load immediate signed doublet big-endian
L.I.16.A.L	Load immediate signed doublet aligned little-endian
L.I.16.L	Load immediate signed doublet little-endian
L.I.32.A.B	Load immediate signed quadlet aligned big-endian
L.I.32.B	Load immediate signed quadlet big-endian
L.I.32.A.L	Load immediate signed quadlet aligned little-endian
L.I.32.L	Load immediate signed quadlet little-endian
L.I.64.A.B	Load immediate signed octlet aligned big-endian
L.I.64.B	Load immediate signed octlet big-endian
L.I.64.A.L	Load immediate signed octlet aligned little-endian
L.I.64.L	Load immediate signed octlet little-endian
L.I.128.A.B ¹⁷	Load immediate hexlet aligned big-endian
L.I.128.B ¹⁸	Load immediate hexlet big-endian
L.I.128.A.L ¹⁹	Load immediate hexlet aligned little-endian
L.I.128.L ²⁰	Load immediate hexlet little-endian
L.I.U.8 ²¹	Load immediate unsigned byte
L.I.U.16.A.B	Load immediate unsigned doublet aligned big-endian
L.I.U.16.B	Load immediate unsigned doublet big-endian
L.I.U.16.A.L	Load immediate unsigned doublet aligned little-endian
L.I.U.16.L	Load immediate unsigned doublet little-endian
L.I.U.32.A.B	Load immediate unsigned quadlet aligned big-endian
L.I.U.32.B	Load immediate unsigned quadlet big-endian
L.I.U.32.A.L	Load immediate unsigned quadlet aligned little-endian
L.I.U.32.L	Load immediate unsigned quadlet little-endian
L.I.U.64.A.B	Load immediate unsigned octlet aligned big-endian
L.I.U.64.B	Load immediate unsigned octlet big-endian
L.I.U.64.A.L	Load immediate unsigned octlet aligned little-endian
L.I.U.64.L	Load immediate unsigned octlet little-endian

¹⁶L.I.8 need not distinguish between little-endian and big-endian ordering, nor between aligned and unaligned, as only a single byte is loaded.

¹⁷L.I.128.AB need not distinguish between signed and unsigned, as the hexlet fills the destination register.

¹⁸L.I.128.B need not distinguish between signed and unsigned, as the hexlet fills the destination register.

¹⁹L.I.128.A.L need not distinguish between signed and unsigned, as the hexlet fills the destination register.

²⁰L.I.128.L need not distinguish between signed and unsigned, as the hexlet fills the destination register.

²¹L.I.U.8 need not distinguish between little-endian and big-endian ordering, nor between aligned and unaligned, as only a single byte is loaded.


Selection

number format	type	size	alignment	ordering
signed byte		8		
unsigned byte	U	8		
signed integer		16 32 64		L B
signed integer aligned		16 32 64	A	L B
unsigned integer	U	16 32 64		L B
unsigned integer aligned	U	16 32 64	A	L B
register		128		L B
register aligned		128	A	L B

Format

op rd=rc,offset

rd=op(rc,offset)

Description

An operand size, expressed in bytes, is specified by the instruction. A virtual address is computed from the sum of the contents of register rc and the sign-extended value of the offset field, multiplied by the operand size. The contents of memory using the specified byte order are read, treated as the size specified, zero-extended or sign-extended as specified, and placed into register rd.

If alignment is specified, the computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an "access disallowed by virtual address" exception occurs.

Definition

```
def LoadImmediate{op,rd,rc,offset} as
  case op of
    LI16L, LI32L, LI8, LI16AL, LI32AL, LI16B, LI32B, LI16AB, LI32AB:
    LI64L, LI64AL, LI64B, LI64AB:
      signed ← true
    LIU16L, LIU32L, LIU8, LIU16AL, LIU32AL,
    LIU16B, LIU32B, LIU16AB, LIU32AB:
    LIU64L, LIU64AL, LIU64B, LIU64AB:
      signed ← false
    LI128L, LI128AL, LI128B, LI128AB:
      signed ← undefined
  endcase
  case op of
    LI8, LIU8:
      size ← 8
    LI16L, LIU16L, LI16AL, LIU16AL, LI16B, LIU16B, LI16AB, LIU16AB:
      size ← 16
```

```

LI32L, LIU32L, LI32AL, LIU32AL, LI32B, LIU32B, LI32AB, LIU32AB:
    size ← 32
LI64L, LIU64L, LI64AL, LIU64AL, LI64B, LIU64B, LI64AB, LIU64AB:
    size ← 64
LI128L, LI128AL, LI128B, LI128AB:
    size ← 128
endcase
lsize ← log(size)
case op of
    LI16L, LIU16L, LI32L, LIU32L, LI64L, LIU64L, LI128L,
    LI16AL, LIU16AL, LI32AL, LIU32AL, LI64AL, LIU64AL, LI128AL:
        order ← LI
    LI16B, LIU16B, LI32B, LIU32B, LI64B, LIU64B, LI128B,
    LI16AB, LIU16AB, LI32AB, LIU32AB, LI64AB, LIU64AB, LI128AB:
        order ← B
    LI8, LIU8:
        order ← undefined
endcase
c ← RegRead[rc, 64]
VirtAddr ← c + {offset55-lsize || offset || 0lsize-31}
case op of
    LI16AL, LIU16AL, LI32AL, LIU32AL, LI64AL, LIU64AL, LI128AL,
    LI16AB, LIU16AB, LI32AB, LIU32AB, LI64AB, LIU64AB, LI128AB:
        if {clsize-4..0 ≠ 0 then
            raise AccessDisallowedByVirtualAddress
        endif
    LI16L, LIU16L, LI32L, LIU32L, LI64L, LIU64L, LI128L,
    LI16B, LIU16B, LI32B, LIU32B, LI64B, LIU64B, LI128B:
    LI8, LIU8:
endcase
m ← LoadMemory[c, VirtAddr, size, order]
a ← {msize-1 and signed}128-size || m
RegWrite[rd, 128, a]
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Store

These operations add the contents of two registers to produce a virtual address, and store the contents of a register into memory.

Operation codes

S.8 ²²	Store byte
S.16.B	Store double big-endian
S.16.AB	Store double aligned big-endian
S.16.L	Store double little-endian
S.16.AL	Store double aligned little-endian
S.32.B	Store quadlet big-endian
S.32.AB	Store quadlet aligned big-endian
S.32.L	Store quadlet little-endian
S.32.AL	Store quadlet aligned little-endian
S.64.B	Store octlet big-endian
S.64.AB	Store octlet aligned big-endian
S.64.L	Store octlet little-endian
S.64.AL	Store octlet aligned little-endian
S.128.B	Store hexlet big-endian
S.128.AB	Store hexlet aligned big-endian
S.128.L	Store hexlet little-endian
S.128.AL	Store hexlet aligned little-endian
S.MUX.64.A.B	Store multiplex octlet aligned big-endian
S.MUX.64.A.L	Store multiplex octlet aligned little-endian


Selection

number format	op	size	alignment	ordering
byte		8		
integer		16 32 64 128		L B
integer aligned		16 32 64 128	A	L B
multiplex	MUX	64	A	L B

Format

op rd,rc,rb

op(rd,rc,rb)



²²S.8 need not specify byte ordering, nor need it specify alignment checking, as it stores a single byte.

Description

An operand size, expressed in bytes, is specified by the instruction. A virtual address is computed from the sum of the contents of register rc and the contents of register rb multiplied by operand size. The contents of register rd, treated as the size specified, is stored in memory using the specified byte order.

If alignment is specified, the computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an "access disallowed by virtual address" exception occurs.

Definition

```
def Store{op,rd,rc,rb} as
  case op of
    S8:
      size ← 8
      S16L, S16AL, S16B, S16AB:
        size ← 16
      S32L, S32AL, S32B, S32AB:
        size ← 32
      S64L, S64AL, S64B, S64AB,
      SMUX64AB, SMUX64AL:
        size ← 64
      S128L, S128AL, S128B, S128AB:
        size ← 128
    endcase
    lsize ← log(size)
  case op of
    S8:
      order ← undefined
      S16L, S32L, S64L, S128L,
      S16AL, S32AL, S64AL, S128AL, SMUX64AL:
        order ← L
      S16B, S32B, S64B, S128B,
      S16AB, S32AB, S64AB, S128AB, SMUX64AB:
        order ← B
    endcase
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 64)
    VirtAddr ← c + (b66-lsize..0 || 0lsize-3)
  case op of
    S16AL, S32AL, S64AL, S128AL,
    S16AB, S32AB, S64AB, S128AB,
    SMUX64AB, SMUX64AL:
      if (clsize-4..0 ≠ 0 then
        raise AccessDisallowedByVirtualAddress
      endif
    S16L, S32L, S64L, S128L,
    S16B, S32B, S64B, S128B:
    S8:
  endcase
  d ← RegRead(rd, 128)
  case op of
```

```

S8,
S16L, S16AL, S16B, S16AB,
S32L, S32AL, S32B, S32AB,
S64L, S64AL, S64B, S64AB,
S128L, S128AL, S128B, S128AB:
    StoreMemory(c,VirtAddr,size,order,dsize-1..0)
SMUX64AB, SMUX64AL:
    lock
        a ← LoadMemoryW(c,VirtAddr,size,order)
        m ← (d127..64 & d63..0) | (a & ~d63..0)
        StoreMemory(c,VirtAddr,size,order,m)
    endlock
endcase
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Store Double Compare Swap

These operations compare two 64-bit values in a register against two 64-bit values read from two 64-bit memory locations, as specified by two 64-bit addresses in a register, and if equal, store two new 64-bit values from a register into the memory locations. The values read from memory are catenated and placed in a register.


Operation codes

S.D.C.S.64.A.B	Store double compare swap octlet aligned big-endian
S.D.C.S.64.A.L	Store double compare swap octlet aligned little-endian

Format

op rd@rc,rb

rd=op(rd,rc,rb)



Description

Two virtual addresses are extracted from the low order bits of the contents of registers rc and rb. Two 64-bit comparison values are extracted from the high order bits of the contents of registers rc and rb. Two 64-bit replacement values are extracted from the contents of register rd. The contents of memory using the specified byte order are read from the specified addresses, treated as 64-bit values, compared against the specified comparison values, and if both read values are equal to the comparison values, the two replacement values are written to memory using the specified byte order. If either are unequal, no values are written to memory. The loaded values are catenated and placed in the register specified by rd.

The virtual addresses must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an “access disallowed by virtual address” exception occurs.

Definition

```
def StoreDoubleCompareSwap(op,rd,rc,rb) as
    size ← 64
    lsize ← log(size)
    case op of
        SDCS64AL:
            order ← L
        SDCS64AB:
```

```

order ← B
endcase
c ← RegRead(rc, 128)
b ← RegRead(rb, 128)
d ← RegRead(rd, 128)
if (c2..0 ≠ 0) or (b2..0 ≠ 0) then
    raise AccessDisallowedByVirtualAddress
endif
lock
    a ← LoadMemoryW(c63..0, c63..0.64, order) || LoadMemoryW(b63..0, b63..0.64, order)
    if ((c127..64 || b127..64) = a) then
        StoreMemory((c63..0, c63..0.64, order, d127..64)
        StoreMemory(b63..0, b63..0.64, order, d63..0)
    endif
endlock
RegWrite(rd, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Store Immediate

These operations add the contents of a register to a sign-extended immediate value to produce a virtual address, and store the contents of a register into memory.

Operation codes

S.I.8 ²³	Store immediate byte
S.I.16.A.B	Store immediate double aligned big-endian
S.I.16.B	Store immediate double big-endian
S.I.16.A.L	Store immediate double aligned little-endian
S.I.16.L	Store immediate double little-endian
S.I.32.A.B	Store immediate quadlet aligned big-endian
S.I.32.B	Store immediate quadlet big-endian
S.I.32.A.L	Store immediate quadlet aligned little-endian
S.I.32.L	Store immediate quadlet little-endian
S.I.64.A.B	Store immediate octlet aligned big-endian
S.I.64.B	Store immediate octlet big-endian
S.I.64.A.L	Store immediate octlet aligned little-endian
S.I.64.L	Store immediate octlet little-endian
S.I.128.A.B	Store immediate hexlet aligned big-endian
S.I.128.B	Store immediate hexlet big-endian
S.I.128.A.L	Store immediate hexlet aligned little-endian
S.I.128.L	Store immediate hexlet little-endian
S.MUXI.64.A.B	Store multiplex immediate octlet aligned big-endian
S.MUXI.64.A.L	Store multiplex immediate octlet aligned little-endian


Selection

number format	op	size	alignment	ordering
byte		8		
integer		16 32 64 128		L B
integer aligned		16 32 64 128	A	L B
multiplex	MUX	64	A	L B

Format

S.op.l.size.align.order rd,rc,offset

sopisizealignorder(rd,rc,offset)



²³S.I.8 need not specify byte ordering, nor need it specify alignment checking, as it stores a single byte.

Description

An operand size, expressed in bytes, is specified by the instruction. A virtual address is computed from the sum of the contents of register rc and the sign-extended value of the offset field, multiplied by the operand size. The contents of register rd, treated as the size specified, are written to memory using the specified byte order.

The computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an “access disallowed by virtual address” exception occurs.

Definition

```
def StoreImmediate(op,rd,rc,offset) as
  case op of
    SI8:
      size ← 8
      SI16L, SI16AL, SI16B, SI16AB:
      size ← 16
      SI32L, SI32AL, SI32B, SI32AB:
      size ← 32
      SI64L, SI64AL, SI64B, SI64AB, SMUXI64AB, SMUXI64AL:
      size ← 64
      SI128L, SI128AL, SI128B, SI128AB:
      size ← 128
    endcase
    lsize ← log(size)
  case op of
    SI8:
      order ← undefined
      SI16L, SI32L, SI64L, SI128L,
      SI16AL, SI32AL, SI64AL, SI128AL, SMUXI64AL:
      order ← L
      SI16B, SI32B, SI64B, SI128B,
      SI16AB, SI32AB, SI64AB, SI128AB, SMUXI64AB:
      order ← B
    endcase
    c ← RegRead(rc, 64)
    VirtAddr ← c + (offset55-lsize || offset || 0lsize-3)
  case op of
    SI16AL, SI32AL, SI64AL, SI128AL,
    SI16AB, SI32AB, SI64AB, SI128AB,
    SMUXI64AB, SMUXI64AL:
      if (lsize-4..0 ≠ 0 then
        raise AccessDisallowedByVirtualAddress
      endif
      SI16L, SI32L, SI64L, SI128L
      SI16B, SI32B, SI64B, SI128B:
    SI8:
  endcase
  d ← RegRead(rd, 128)
  case op of
    SI8,
    SI16L, SI16AL, SI16B, SI16AB,
    SI32L, SI32AL, SI32B, SI32AB,
```

```
SI64L, SI64AL, SI64B, SI64AB,  
SI128L, SI128AL, SI128B, SI128AB:  
    StoreMemory(c,VirtAddr,size,order,dsize-1..0)  
SMUXI64AB, SMUXI64AL:  
    lock  
        a ← LoadMemoryW(c,VirtAddr,size,order)  
        m ← (d127..64 & d63..0) | (a & ~d63..0)  
        StoreMemory(c,VirtAddr,size,order,m)  
    endlock  
endcase  
enddef
```

Exceptions

Access disallowed by virtual address
Access disallowed by tag
Access disallowed by global TB
Access disallowed by local TB
Access detail required by tag
Access detail required by local TB
Access detail required by global TB
Local TB miss
Global TB miss

Store Immediate Inplace

These operations add the contents of a register to a sign-extended immediate value to produce a virtual address, and store the contents of a register into memory.

Operation codes

S.A.S.I.64.A.B	Store add swap immediate octlet aligned big-endian
S.A.S.I.64.A.L	Store add swap immediate octlet aligned little-endian
S.C.S.I.64.A.B	Store compare swap immediate octlet aligned big-endian
S.C.S.I.64.A.L	Store compare swap immediate octlet aligned little-endian
S.M.S.I.64.A.B	Store multiplex swap immediate octlet aligned big-endian
S.M.S.I.64.A.L	Store multiplex swap immediate octlet aligned little-endian


Selection

number format	op	size	alignment	ordering
add-swap	AS	64	A	L B
compare-swap	CS	64	A	L B
multiplex-swap	MS	64	A	L B

Format

S.op.I.64.align.order rd@rc,offset

rd=sopi64alignorder(rd,rc,offset)



Description

A virtual address is computed from the sum of the contents of register *rc* and the sign-extended value of the *offset* field. The contents of memory using the specified byte order are read and treated as a 64-bit value. A specified operation is performed between the memory contents and the original contents of register *rd*, and the result is written to memory using the specified byte order. The original memory contents are placed into register *rd*.

The computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an “access disallowed by virtual address” exception occurs.

Definition

```
def StoreImmediateInplace(op,rd,rc,offset) as
    size ← 64
    lsize ← log(size)
    case op of
```

```

SASI64AL, SCSI64AL, SMSI64AL:
    order ← L
SASI64AB, SCSI64AB, SMSI64AB:
    order ← B
endcase
c ← RegRead(rc, 64)
VirtAddr ← c + {offset55-lsize || offset || 0lsize-3}
if (Clsize-4..0 ≠ 0 then
    raise AccessDisallowedByVirtualAddress
endif
d ← RegRead(rd, 128)
case op of
    SASI64AB, SASI64AL:
        lock
            a ← LoadMemoryW(c,VirtAddr,size,order)
            StoreMemory(c,VirtAddr,size,order,d63..0+a)
        endlock
    SCSI64AB, SCSI64AL:
        lock
            a ← LoadMemoryW(c,VirtAddr,size,order)
            if (a = d63..0) then
                StoreMemory(c,VirtAddr,size,order,d127..64)
            endif
        endlock
    SMSI64AB, SMSI64AL:
        lock
            a ← LoadMemoryW(c,VirtAddr,size,order)
            m ← (d127..64 & d63..0) | (a & -d63..0)
            StoreMemory(c,VirtAddr,size,order,m)
        endlock
endcase
RegWrite(rd, 64, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Store Inplace

These operations add the contents of two registers to produce a virtual address, and store the contents of a register into memory.

Operation codes

S.A.S.64.A.B	Store add swap octlet aligned big-endian
S.A.S.64.A.L	Store add swap octlet aligned little-endian
S.C.S.64.A.B	Store compare swap octlet aligned big-endian
S.C.S.64.A.L	Store compare swap octlet aligned little-endian
S.M.S.64.A.B	Store multiplex swap octlet aligned big-endian
S.M.S.64.A.L	Store multiplex swap octlet aligned little-endian


Selection

number format	op	size	alignment	ordering
add-swap	A.S	64	A	L B
compare-swap	C.S	64	A	L B
multiplex-swap	M.S	64	A	L B

Format

op rd@rc,rb

rd=op(rd,rc,rb)



Description

A virtual address is computed from the sum of the contents of register rc and the contents of register rb multiplied by operand size. The contents of memory using the specified byte order are read and treated as 64 bits. A specified operation is performed between the memory contents and the original contents of register rd, and the result is written to memory using the specified byte order. The original memory contents are placed into register rd.

The computed virtual address must be aligned, that is, it must be an exact multiple of the size expressed in bytes. If the address is not aligned an “access disallowed by virtual address” exception occurs.

Definition

```
def StoreInplace(op,rd,rc,rb) as
    size ← 64
    lsize ← log(size)
    case op of
```

```

SAS64AL, SCS64AL, SMS64AL:
    order ← L
SAS64AB, SCS64AB, SMS64AB:
    order ← B
endcase
c ← RegRead[rc, 64]
b ← RegRead[rb, 64]
VirtAddr ← c + {b66-lsize..0 || 0lsize-3}
if (lsize-4..0 ≠ 0 then
    raise AccessDisallowedByVirtualAddress
endif
d ← RegRead[rd, 128]
case op of
    SAS64AB, SAS64AL:
        lock
            a ← LoadMemoryW(c,VirtAddr,size,order)
            StoreMemory(c,VirtAddr,size,order,d63..0+a)
        endlock
    SCS64AB, SCS64AL:
        lock
            a ← LoadMemoryW(c,VirtAddr,size,order)
            if (a = d63..0) then
                StoreMemory(c,VirtAddr,size,order,d127..64)
            endif
        endlock
    SMS64AB, SMS64AL:
        lock
            a ← LoadMemoryW(c,VirtAddr,size,order)
            m ← {d127..64 & d63..0} | (a & ~d63..0)
            StoreMemory(c,VirtAddr,size,order,m)
        endlock
endcase
RegWrite[rd, 64, a]
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Group Add

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

G.ADD.8	Group add bytes
G.ADD.16	Group add doublets
G.ADD.32	Group add quadlets
G.ADD.64	Group add octlets
G.ADD.128	Group add hexlet
GADD.L.8	Group add limit signed bytes
GADD.L.16	Group add limit signed doublets
GADD.L.32	Group add limit signed quadlets
GADD.L.64	Group add limit signed octlets
GADD.L.128	Group add limit signed hexlet
GADD.L.U.8	Group add limit unsigned bytes
GADD.L.U.16	Group add limit unsigned doublets
GADD.L.U.32	Group add limit unsigned quadlets
GADD.L.U.64	Group add limit unsigned octlets
GADD.L.U.128	Group add limit unsigned hexlet
G.ADD.8.O	Group add signed bytes check overflow
G.ADD.16.O	Group add signed doublets check overflow
G.ADD.32.O	Group add signed quadlets check overflow
G.ADD.64.O	Group add signed octlets check overflow
G.ADD.128.O	Group add signed hexlet check overflow
G.ADD.U.8.O	Group add unsigned bytes check overflow
G.ADD.U.16.O	Group add unsigned doublets check overflow
G.ADD.U.32.O	Group add unsigned quadlets check overflow
G.ADD.U.64.O	Group add unsigned octlets check overflow
G.ADD.U.128.O	Group add unsigned hexlet check overflow


Redundancies

GADD.size rd=rc,rc	\Leftrightarrow	G.SHL.I.size rd=rc,1
GADD.size.O rd=rc,rc	\Leftrightarrow	G.SHL.I.size.O rd=rc,1
G.ADD.U.size.O rd=rc,rc	\Leftrightarrow	G.SHL.I.U.size.O rd=rc,1

Format

G.op.sizerd=rc,rb

rd=gopsize(rc,rb)

Description

The contents of registers rc and rb are partitioned into groups of the size specified and added, and if specified, checked for overflow or limited, yielding a group of results, each of which is the size specified. The group of results is catenated and placed in register rd.

Definition

```
def Group{op,size,rd,rc,rb}
    c ← RegRead[rc, 128]
    b ← RegRead[rb, 128]
    case op of
        G.ADD:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ci+size-1..i + bi+size-1..i
            endfor
        G.ADD.L:
            for i ← 0 to 128-size by size
                t ← (ci+size-1 || ci+size-1..i) + (bi+size-1 || bi+size-1..i)
                ai+size-1..i ← (tsize ≠ tsize-1) ? (tsize || tsize-1) : tsize-1..0
            endfor
        G.ADD.LU:
            for i ← 0 to 128-size by size
                t ← (0^ || ci+size-1..i) + (0^ || bi+size-1..i)
                ai+size-1..i ← (tsize ≠ 0) ? (1^size) : tsize-1..0
            endfor
        G.ADD.O:
            for i ← 0 to 128-size by size
                t ← (ci+size-1 || ci+size-1..i) + (bi+size-1 || bi+size-1..i)
                if tsize ≠ tsize-1 then
                    raise FixedPointArithmetic
                endif
                ai+size-1..i ← tsize-1..0
            endfor
        G.ADD.UO:
            for i ← 0 to 128-size by size
                t ← (0^ || ci+size-1..i) + (0^ || bi+size-1..i)
                if tsize ≠ 0 then
                    raise FixedPointArithmetic
                endif
                ai+size-1..i ← tsize-1..0
            endfor
```

```
    endfor
  endcase
  RegWrite(rd, 128, a)
enddef
```

Exceptions

Fixed-point arithmetic

Group Add Halve

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

GADD.H.8.C	Group add halve signed bytes ceiling
GADD.H.8.F	Group add halve signed bytes floor
GADD.H.8.N	Group add halve signed bytes nearest
GADD.H.8.Z	Group add halve signed bytes zero
GADD.H.16.C	Group add halve signed doublets ceiling
GADD.H.16.F	Group add halve signed doublets floor
GADD.H.16.N	Group add halve signed doublets nearest
GADD.H.16.Z	Group add halve signed doublets zero
GADD.H.32.C	Group add halve signed quadlets ceiling
GADD.H.32.F	Group add halve signed quadlets floor
GADD.H.32.N	Group add halve signed quadlets nearest
GADD.H.32.Z	Group add halve signed quadlets zero
GADD.H.64.C	Group add halve signed octlets ceiling
GADD.H.64.F	Group add halve signed octlets floor
GADD.H.64.N	Group add halve signed octlets nearest
GADD.H.64.Z	Group add halve signed octlets zero
GADD.H.128.C	Group add halve signed hexlet ceiling
GADD.H.128.F	Group add halve signed hexlet floor
GADD.H.128.N	Group add halve signed hexlet nearest
GADD.H.128.Z	Group add halve signed hexlet zero
GADD.H.U.8.C	Group add halve unsigned bytes ceiling
GADD.H.U.8.F	Group add halve unsigned bytes floor
GADD.H.U.8.N	Group add halve unsigned bytes nearest
GADD.H.U.16.C	Group add halve unsigned doublets ceiling
GADD.H.U.16.F	Group add halve unsigned doublets floor
GADD.H.U.16.N	Group add halve unsigned doublets nearest
GADD.H.U.32.C	Group add halve unsigned quadlets ceiling
GADD.H.U.32.F	Group add halve unsigned quadlets floor
GADD.H.U.32.N	Group add halve unsigned quadlets nearest
GADD.H.U.64.C	Group add halve unsigned octlets ceiling
GADD.H.U.64.F	Group add halve unsigned octlets floor
GADD.H.U.64.N	Group add halve unsigned octlets nearest
GADD.H.U.128.C	Group add halve unsigned hexlet ceiling
GADD.H.U.128.F	Group add halve unsigned hexlet floor
GADD.H.U.128.N	Group add halve unsigned hexlet nearest


Redundancies

GADD.H.size.rnd rd=rc,rc	\Leftrightarrow	G.COPY rd=rc
GADD.H.U.size.rnd rd=rc,rc	\Leftrightarrow	G.COPY rd=rc

Format

G.op.size.rnd rd=rc,rb

rd=gopsizernd(rc,rb)

Description

The contents of registers rc and rb are partitioned into groups of operands of the size specified, added, halved, and rounded as specified, yielding a group of results, each of which is the size specified. The results never overflow, so limiting is not required by this operation. The group of results is catenated and placed in register rd.

Z (zero) rounding is not defined for unsigned operations, and a ReservedInstruction exception is raised if attempted. F (floor) rounding will properly round unsigned results downward.

Definition

```
def GroupAddHalve(op,rnd,size,rd,rc,rb)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case op of
        GADDHC, GADDHF, GADDHN, GADDHZ:
            as ← cs ← bs ← 1
        GADDHUC, GADDHUF, GADDHUN, GADDHUZ
            as ← cs ← bs ← 0
            if rnd = Z then
                raise ReservedInstruction
            endif
    endcase
    h ← size+1
    r ← 1
    for i ← 0 to 128-size by size
        p ← ((cs and csize-1) || csize-1+i..i) + ((bs and bsize-1) || bsize-1+i..i)
        case rnd of
            none, N:
                s ← 0size || -p;
            Z:
                s ← 0size || psize
            F:
                s ← 0size+1
            C:
```

```
s ← 0size || 11
endcase
v ← {(as & psize) || p} + (0 || s)
a[size-1+i..i] ← v[size..1]
endfor
RegWrite(rd, 128, a)
enddef
```

Exceptions

ReservedInstruction

Group Boolean

These operations take operands from three registers, perform boolean operations on corresponding bits in the operands, and place the concatenated results in the third register.

Operation codes

G.BOOLEAN	Group boolean
-----------	---------------

Equivalencies

G.AAA	Group three-way and
G.AAA.1	Group add add add bits
G.AAS.1	Group add add subtract bits
G.ADD.1	Group add bits
G.AND	Group and
G.ANDN	Group and not
G.COPY	Group copy
G.NAAA	Group three-way nand
G.NAND	Group nand
G.NOOO	Group three-way nor
G.NOR	Group nor
G.NOT	Group not
G.NXXX	Group three-way exclusive-nor
G.OOO	Group three-way or
G.OR	Group or
G.ORN	Group or not
G.SAA.1	Group subtract add add bits
G.SAS.1	Group subtract add subtract bits
G.SET	Group set
G.SET.AND.E.1	Group set and equal zero bits
G.SET.AND.NE.1	Group set and not equal zero bits
G.SET.E.1	Group set equal bits
G.SET.G.1	Group set greater signed bits
G.SET.G.U.1	Group set greater unsigned bits
G.SET.G.Z.1	Group set greater zero signed bits
G.SET.GE.1	Group set greater equal signed bits
G.SET.GE.Z.1	Group set greater equal zero signed bits
G.SET.L.1	Group set less signed bits
G.SET.L.Z.1	Group set less zero signed bits
G.SET.LE.1	Group set less equal signed bits
G.SET.LE.U.1	Group set less equal unsigned bits
G.SET.LE.Z.1	Group set less equal zero signed bits
G.SET.NE.1	Group set not equal bits
G.SET.GE.U.1	Group set greater equal unsigned bits
G.SET.L.U.1	Group set less unsigned bits

G.SSA.1	Group subtract subtract add bits
G.SSS.1	Group subtract subtract subtract bits
G.SUB.1	Group subtract bits
G.XNOR	Group exclusive-nor
G.XOR	Group exclusive-or
G.XXX	Group three-way exclusive-or
G.ZERO	Group zero

G.AAA rd@rc,rb	← G.BOOLEAN rd@rc,rb,0b10000000
G.AAA.1 rd@rc,rb	→ G.XXX rd@rc,rb
G.AAS.1 rd@rc,rb	→ G.XXX rd@rc,rb
G.ADD.1 rd=rc,rb	→ G.XOR rd=rc,rb
G.AND rd=rc,rb	← G.BOOLEAN rd@rc,rb,0b10001000
G.ANDN rd=rc,rb	← G.BOOLEAN rd@rc,rb,0b01000100
G.BOOLEAN rd@rb,rc,i	→ G.BOOLEAN rd@rc,rb,i ₇ i ₅ i ₆ i ₄ i ₃ i ₁ i ₂ i ₀
G.COPY rd=rc	← G.BOOLEAN rd@rc,rc,0b10001000
G.NAAA rd@rc,rb	← G.BOOLEAN rd@rc,rb,0b01111111
G.NAND rd=rc,rb	← G.BOOLEAN rd@rc,rb,0b01110111
G.NOOO rd@rc,rb	← G.BOOLEAN rd@rc,rb,0b00000001
G.NOR rd=rc,rb	← G.BOOLEAN rd@rc,rb,0b00010001
G.NOT rd=rc	← G.BOOLEAN rd@rc,rc,0b00010001
G.NXXX rd@rc,rb	← G.BOOLEAN rd@rc,rb,0b01101001
G.OOO rd@rc,rb	← G.BOOLEAN rd@rc,rb,0b11111110
G.OR rd=rc,rb	← G.BOOLEAN rd@rc,rb,0b11101110
G.ORN rd=rc,rb	← G.BOOLEAN rd@rc,rb,0b11011101
G.SAA.1 rd@rc,rb	→ G.XXX rd@rc,rb
G.SAS.1 rd@rc,rb	→ G.XXX rd@rc,rb
G.SET rd	← G.BOOLEAN rd@rd,rd,0b10000001
G.SET.AND.E.1 rd=rb,rc	→ G.NAND rd=rc,rb
G.SET.AND.NE.1 rd=rb,rc	→ G.AND rd=rc,rb
G.SET.E.1 rd=rb,rc	→ G.XNOR rd=rc,rb
G.SET.G.1 rd=rb,rc	→ G.ANDN rd=rc,rb
G.SET.G.U.1 rd=rb,rc	→ G.ANDN rd=rb,rc
G.SET.G.Z.1 rd=rc	→ G.ZERO rd
G.SET.GE.1 rd=rb,rc	→ G.ORN rd=rc,rb
G.SET.GE.Z.1 rd=rc	→ G.NOT rd=rc
G.SET.L.1 rd=rb,rc	→ G.ANDN rd=rb,rc
G.SET.L.Z.1 rd=rc	→ G.COPY rd=rc
G.SET.LE.1 rd=rb,rc	→ G.ORN rd=rb,rc
G.SET.LE.U.1 rd=rb,rc	→ G.ORN rd=rc,rb
G.SET.LE.Z.1 rd=rc	→ G.SET rd
G.SET.NE.1 rd=rb,rc	→ G.XOR rd=rc,rb

<i>G.SET.GE.U.1 rd=rb,rc</i>	→ G.ORN rd=rb,rc
<i>G.SET.L.U.1 rd=rb,rc</i>	→ GANDN rd=rc,rb
<i>G.SSA.1 rd@rc,rb</i>	→ GXXX rd@rc,rb
<i>G.SSS.1 rd@rc,rb</i>	→ GXXX rd@rc,rb
<i>G.SUB.1 rd=rc,rb</i>	→ GXOR rd=rc,rb
<i>G.XNOR rd=rc,rb</i>	← G.BOOLEAN rd@rc,rb,0b10011001
<i>G.XOR rd=rc,rb</i>	← G.BOOLEAN rd@rc,rb,0b01100110
<i>G.XXX rd@rc,rb</i>	← G.BOOLEAN rd@rc,rb,0b10010110
<i>G.ZERO rd</i>	← G.BOOLEAN rd@rd,rd,0b00000000


Selection

operation	function (binary)	function (decimal)
d	11110000	240
c	11001100	204
b	10101010	176
d&c&b	10000000	128
(d&c)1b	11101010	234
d1c1b	11111110	254
d?c:b	11001010	202
d^c^b	10010110	150
-d^c^b	01101001	105
0	00000000	0

Format

G.BOOLEAN rd@trc,trb,f

rd=gbooleani(rd,rc,rb,f)



```

if f6=f5 then
    if f2=f1 then
        if f2 then
            rc ← max(trc,trb)
            rb ← min(trc,trb)
        else
            rc ← min(trc,trb)
            rb ← max(trc,trb)
        endif
        ih ← 0
        il ← 0 || f6 || f7 || f4 || f3 || f0
    else
        if f2 then
            rc ← trb
            rb ← trc
        else
            rc ← trc
            rb ← trb
        endif
        ih ← 0
        il ← 1 || f6 || f7 || f4 || f3 || f0
    endif
else
    ih ← 1
    if f6 then
        rc ← trb
        rb ← trc
        il ← f1 || f2 || f7 || f4 || f3 || f0
    else
        rc ← trc
        rb ← trb
        il ← f2 || f1 || f7 || f4 || f3 || f0
    endif
endif

```

Description

Three values are taken from the contents of registers rd, rc and rb. The ih and il fields specify a function of three bits, producing a single bit result. The specified function is evaluated for each bit position, and the results are catenated and placed in register rd.

Register rd is both a source and destination of this instruction.

The function is specified by eight bits, which give the result for each possible value of the three source bits in each bit position:

d	1 1 1 1 0 0 0 0
c	1 1 0 0 1 1 0 0
b	1 0 1 0 1 0 1 0
f(d,c,b)	f ₇ f ₆ f ₅ f ₄ f ₃ f ₂ f ₁ f ₀

A function can be modified by rearranging the bits of the immediate value. The table below shows how rearrangement of immediate value f_{7..0} can reorder the operands d,c,b for the same function.

operation	immediate
f(d,c,b)	f ₇ f ₆ f ₅ f ₄ f ₃ f ₂ f ₁ f ₀
f(c,d,b)	f ₇ f ₆ f ₃ f ₂ f ₅ f ₄ f ₁ f ₀
f(d,b,c)	f ₇ f ₅ f ₆ f ₄ f ₃ f ₁ f ₂ f ₀
f(b,c,d)	f ₇ f ₃ f ₅ f ₁ f ₆ f ₂ f ₄ f ₀
f(c,b,d)	f ₇ f ₅ f ₃ f ₁ f ₆ f ₄ f ₂ f ₀
f(b,d,c)	f ₇ f ₃ f ₆ f ₂ f ₅ f ₁ f ₄ f ₀

By using such a rearrangement, an operation of the form: $b=f(d,c,b)$ can be recoded into a legal form: $b=f(b,d,c)$. For example, the function: $b=f(d,c,b)=d?c:b$ cannot be coded, but the equivalent function: $d=c?b:d$ can be determined by rearranging the code for $d=f(d,c,b)=d?c:b$, which is 11001010, according to the rule for $f(d,c,b) \Rightarrow f(c,b,d)$, to the code 11011000.

Encoding

Some special characteristics of this rearrangement is the basis of the manner in which the eight function specification bits are compressed to seven immediate bits in this instruction. As seen in the table above, in the general case, a rearrangement of operands from $f(d,c,b)$ to $f(d,b,c)$ (interchanging rc and rb) requires interchanging the values of f₆ and f₅ and the values of f₂ and f₁.

Among the 256 possible functions which this instruction can perform, one quarter of them (64 functions) are unchanged by this rearrangement. These functions have the property that f₆=f₅ and f₂=f₁. The values of rc and rb²⁴ can be freely interchanged, and so are sorted into rising or falling order to indicate the value of f₂.²⁵ These functions are encoded by the values of f₇, f₆, f₄, f₃, and f₀ in the immediate field and f₂ by whether rc>rb, thus using 32 immediate values for 64 functions.

Another quarter of the functions have f₆=1 and f₅=0. These functions are recoded by interchanging rc and rb, f₆ and f₅, f₂ and f₁. They then share the same encoding as the

²⁴ Note that rc and rb are the register specifiers, not the register contents.

²⁵ A special case arises when rc=rb, so the sorting of rc and rb cannot convey information. However, as only the values f₇, f₄, f₃, and f₀ can ever result in this case, f₆, f₅, f₂, and f₁ need not be coded for this case, so no special handling is required.

quarter of the functions where $f_6=0$ and $f_5=1$, and are encoded by the values of f_7, f_4, f_3, f_2, f_1 , and f_0 in the immediate field, thus using 64 immediate values for 128 functions.

The remaining quarter of the functions have $f_6=f_5$ and $f_2 \neq f_1$. The half of these in which $f_2=1$ and $f_1=0$ are recoded by interchanging rc and rb, f_6 and f_5 , f_2 and f_1 . They then share the same encoding as the eighth of the functions where $f_2=0$ and $f_1=1$, and are encoded by the values of f_7, f_6, f_4, f_3 , and f_0 in the immediate field, thus using 32 immediate values for 64 functions.

The function encoding is summarized by the table:

f_7	f_6	f_5	f_4	f_3	f_2	f_1	f_0	$trc > trb$	ih	il_5	il_4	il_3	il_2	il_1	il_0	rc	rb
	f_6				f_2		f_2		0	0	f_6	f_7	f_4	f_3	f_0	trc	trb
	f_6				f_2		$\neg f_2$		0	0	f_6	f_7	f_4	f_3	f_0	trb	trc
	f_6			0	1				0	1	f_6	f_7	f_4	f_3	f_0	trc	trb
	f_6			1	0				0	1	f_6	f_7	f_4	f_3	f_0	trb	trc
0	1								1	f_2	f_1	f_7	f_4	f_3	f_0	trc	trb
1	0								1	f_1	f_2	f_7	f_4	f_3	f_0	trb	trc

The function decoding is summarized by the table:

ih	il_5	il_4	il_3	il_2	il_1	il_0	$rc > rb$	f_7	f_6	f_5	f_4	f_3	f_2	f_1	f_0
0	0						0	il_3	il_4	il_4	il_2	il_1	0	0	il_0
0	0						1	il_3	il_4	il_4	il_2	il_1	1	1	il_0
0	1							il_3	il_4	il_4	il_2	il_1	0	1	il_0
1								il_3	0	1	il_2	il_1	il_5	il_4	il_0

Definition

```

def GroupBoolean {ih,rd,rc,rb,il}
  d ← RegRead(rd, 128)
  c ← RegRead(rc, 128)
  b ← RegRead(rb, 128)
  if  $ih=0$  then
    if  $il_5=0$  then
      f ←  $il_3 \parallel il_4 \parallel il_4 \parallel il_2 \parallel il_1 \parallel (rc > rb)^2 \parallel il_0$ 
    else
      f ←  $il_3 \parallel il_4 \parallel il_4 \parallel il_2 \parallel il_1 \parallel 0 \parallel 1 \parallel 1 \parallel il_0$ 
    endif
  else
    f ←  $il_3 \parallel 0 \parallel 1 \parallel 1 \parallel il_2 \parallel il_1 \parallel il_5 \parallel il_4 \parallel il_0$ 
  endif
  for i ← 0 to 127 by size
    ai ← fdi || ci || bi
  endfor
  RegWrite(rd, 128, a)
enddef

```

Exceptions

none

Group Compare

These operations perform calculations on partitions of bits in two general register values, and generate a fixed-point arithmetic exception if the condition specified is met.

Operation codes

G.COM.AND.E.8	Group compare and equal zero bytes
G.COM.AND.E.16	Group compare and equal zero doublets
G.COM.AND.E.32	Group compare and equal zero quadlets
G.COM.AND.E.64	Group compare and equal zero octlets
G.COM.AND.E.128	Group compare and equal zero hexlet
G.COM.AND.NE.8	Group compare and not equal zero bytes
G.COM.AND.NE.16	Group compare and not equal zero doublets
G.COM.AND.NE.32	Group compare and not equal zero quadlets
G.COM.AND.NE.64	Group compare and not equal zero octlets
G.COM.AND.NE.128	Group compare and not equal zero hexlet
G.COM.E.8	Group compare equal bytes
G.COM.E.16	Group compare equal doublets
G.COM.E.32	Group compare equal quadlets
G.COM.E.64	Group compare equal octlets
G.COM.E.128	Group compare equal hexlet
G.COM.GE.8	Group compare greater equal signed bytes
G.COM.GE.16	Group compare greater equal signed doublets
G.COM.GE.32	Group compare greater equal signed quadlets
G.COM.GE.64	Group compare greater equal signed octlets
G.COM.GE.128	Group compare greater equal signed hexlet
G.COM.GE.U.8	Group compare greater equal unsigned bytes
G.COM.GE.U.16	Group compare greater equal unsigned doublets
G.COM.GE.U.32	Group compare greater equal unsigned quadlets
G.COM.GE.U.64	Group compare greater equal unsigned octlets
G.COM.GE.U.128	Group compare greater equal unsigned hexlet
G.COM.L.8	Group compare signed less bytes
G.COM.L.16	Group compare signed less doublets
G.COM.L.32	Group compare signed less quadlets
G.COM.L.64	Group compare signed less octlets
G.COM.L.128	Group compare signed less hexlet
G.COM.L.U.8	Group compare less unsigned bytes
G.COM.L.U.16	Group compare less unsigned doublets
G.COM.L.U.32	Group compare less unsigned quadlets
G.COM.L.U.64	Group compare less unsigned octlets
G.COM.L.U.128	Group compare less unsigned hexlet
G.COM.NE.8	Group compare not equal bytes
G.COM.NE.16	Group compare not equal doublets
G.COM.NE.32	Group compare not equal quadlets
G.COM.NE.64	Group compare not equal octlets

G.COM.NE.128

Group compare not equal hexlet

Equivalencies

G.COM.E.Z.8	Group compare equal zero signed bytes
G.COM.E.Z.16	Group compare equal zero signed doublets
G.COM.E.Z.32	Group compare equal zero signed quadlets
G.COM.E.Z.64	Group compare equal zero signed octlets
G.COM.E.Z.128	Group compare equal zero signed hexlet
G.COM.G.8	Group compare signed greater bytes
G.COM.G.16	Group compare signed greater doublets
G.COM.G.32	Group compare signed greater quadlets
G.COM.G.64	Group compare signed greater octlets
G.COM.G.128	Group compare signed greater hexlet
G.COM.G.U.8	Group compare greater unsigned bytes
G.COM.G.U.16	Group compare greater unsigned doublets
G.COM.G.U.32	Group compare greater unsigned quadlets
G.COM.G.U.64	Group compare greater unsigned octlets
G.COM.G.U.128	Group compare greater unsigned hexlet
G.COM.G.Z.8	Group compare greater zero signed bytes
G.COM.G.Z.16	Group compare greater zero signed doublets
G.COM.G.Z.32	Group compare greater zero signed quadlets
G.COM.G.Z.64	Group compare greater zero signed octlets
G.COM.G.Z.128	Group compare greater zero signed hexlet
G.COM.GE.Z.8	Group compare greater equal zero signed bytes
G.COM.GE.Z.16	Group compare greater equal zero signed doublets
G.COM.GE.Z.32	Group compare greater equal zero signed quadlets
G.COM.GE.Z.64	Group compare greater equal zero signed octlets
G.COM.GE.Z.128	Group compare greater equal zero signed hexlet
G.COM.L.Z.8	Group compare less zero signed bytes
G.COM.L.Z.16	Group compare less zero signed doublets
G.COM.L.Z.32	Group compare less zero signed quadlets
G.COM.L.Z.64	Group compare less zero signed octlets
G.COM.L.Z.128	Group compare less zero signed hexlet
G.COM.LE.8	Group compare less equal signed bytes
G.COM.LE.16	Group compare less equal signed doublets
G.COM.LE.32	Group compare less equal signed quadlets
G.COM.LE.64	Group compare less equal signed octlets
G.COM.LE.128	Group compare less equal signed hexlet
G.COM.LE.U.8	Group compare less equal unsigned bytes
G.COM.LE.U.16	Group compare less equal unsigned doublets
G.COM.LE.U.32	Group compare less equal unsigned quadlets
G.COM.LE.U.64	Group compare less equal unsigned octlets
G.COM.LE.U.128	Group compare less equal unsigned hexlet
G.COM.LE.Z.8	Group compare less equal zero signed bytes
G.COM.LE.Z.16	Group compare less equal zero signed doublets
G.COM.LE.Z.32	Group compare less equal zero signed quadlets
G.COM.LE.Z.64	Group compare less equal zero signed octlets

G.COM.LE.Z.128	Group compare less equal zero signed hexlet
G.COM.NE.Z.8	Group compare not equal zero signed bytes
G.COM.NE.Z.16	Group compare not equal zero signed doublets
G.COM.NE.Z.32	Group compare not equal zero signed quadlets
G.COM.NE.Z.64	Group compare not equal zero signed octlets
G.COM.NE.Z.128	Group compare not equal zero signed hexlet
G.FIX	Group fixed point arithmetic exception
G.NOP	Group no operation

G.COM.E.Z.size rc	← G.COMAND.E.size rc,rc
G.COM.G.size rd,rc	→ G.COM.L.size rc,rd
G.COM.G.U.size rd,rc	→ G.COM.L.U.size rc,rd
G.COM.G.Z.size rc	← G.COM.L.U.size rc,rc
G.COM.GE.Z.size rc	← G.COM.GE.size rc,rc
G.COM.L.Z.size rc	← G.COM.L.size rc,rc
G.COM.LE.size rd,rc	→ G.COM.GE.size rc,rd
G.COM.LE.U.size rd,rc	→ G.COM.GE.U.size rc,rd
G.COM.LE.Z.size rc	← G.COM.GE.U.size rc,rc
G.COM.NE.Z.size rc	← G.COMAND.NE.size rc,rc
G.FIX	← G.COM.E.128 r0,r0
G.NOP	← G.COM.NE.128 r0,r0

Redundancies

G.COM.E.size rd,rd	↔ G.FIX
G.COM.NE.size rd,rd	↔ G.NOP


Selection

class	operation	cond	type	size
boolean	COMAND	E NE	NONE U	8 16 32 64 128
	COM			
arithmetic	COM	L GE G LE	Z	8 16 32 64 128
	COM	L GE G LE E NE		8 16 32 64 128

Format

G.COM.op.size rd,rc
 G.COM.opz.size rcd

gcomopsizex(rd,rc)



Description

Two values are taken from the contents of registers rd and rc. The specified condition is calculated on partitions of the operands. If the specified condition is true for any partition, a fixed-point arithmetic exception is generated. This instruction generates no general purpose register results..

Definition

```
def GroupCompare(op,size,rd,rc)
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    case op of
        G.COM.E:
            for i ← 0 to 128-size by size
                ai+size-1..i ← (di+size-1..i = ci+size-1..i)size
            endfor
        G.COM.NE:
            for i ← 0 to 128-size by size
                ai+size-1..i ← (di+size-1..i ≠ ci+size-1..i)size
            endfor
        G.COM.AND.E:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ((ci+size-1..i and di+size-1..i) = 0)size
            endfor
        G.COM.AND.NE:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ((ci+size-1..i and di+size-1..i) ≠ 0)size
            endfor
        G.COM.I.:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ((rd = rc) ? (ci+size-1..i < 0) : (di+size-1..i < ci+size-1..i))size
            endfor
        G.COM.GE.:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ((rd = rc) ? (ci+size-1..i ≥ 0) : (di+size-1..i ≥ ci+size-1..i))size
            endfor
        G.COM.LU.:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ((rd = rc) ? (ci+size-1..i > 0) :
                    ((0 || di+size-1..i) < (0 || ci+size-1..i)))size
            endfor
        G.COM.GE.U.:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ((rd = rc) ? (ci+size-1..i ≤ 0) :
                    ((0 || di+size-1..i) ≥ (0 || ci+size-1..i)))size
            endfor
    endcase
    if (a ≠ 0) then
        raise FixedPointArithmetic
    endif
enddef
```

Exceptions

Fixed-point arithmetic

Group Compare Floating-point

These operations perform calculations on partitions of bits in two general register values, and generate a floating-point arithmetic exception if the condition specified is met.

Operation codes

G.COM.E.F.16	Group compare equal floating-point half
G.COM.E.F.16.X	Group compare equal floating-point half exact
G.COM.E.F.32	Group compare equal floating-point single
G.COM.E.F.32.X	Group compare equal floating-point single exact
G.COM.E.F.64	Group compare equal floating-point double
G.COM.E.F.64.X	Group compare equal floating-point double exact
G.COM.E.F.128	Group compare equal floating-point quad
G.COM.E.F.128.X	Group compare equal floating-point quad exact
G.COM.GE.F.16	Group compare greater or equal floating-point half
G.COM.GE.F.16.X	Group compare greater or equal floating-point half exact
G.COM.GE.F.32	Group compare greater or equal floating-point single
G.COM.GE.F.32.X	Group compare greater or equal floating-point single exact
G.COM.GE.F.64	Group compare greater or equal floating-point double
G.COM.GE.F.64.X	Group compare greater or equal floating-point double exact
G.COM.GE.F.128	Group compare greater or equal floating-point quad
G.COM.GE.F.128.X	Group compare greater or equal floating-point quad exact
G.COM.L.F.16	Group compare less floating-point half
G.COM.L.F.16.X	Group compare less floating-point half exact
G.COM.L.F.32	Group compare less floating-point single
G.COM.L.F.32.X	Group compare less floating-point single exact
G.COM.L.F.64	Group compare less floating-point double
G.COM.L.F.64.X	Group compare less floating-point double exact
G.COM.L.F.128	Group compare less floating-point quad
G.COM.L.F.128.X	Group compare less floating-point quad exact
G.COM.LG.F.16	Group compare less or greater floating-point half
G.COM.LG.F.16.X	Group compare less or greater floating-point half exact
G.COM.LG.F.32	Group compare less or greater floating-point single
G.COM.LG.F.32.X	Group compare less or greater floating-point single exact
G.COM.LG.F.64	Group compare less or greater floating-point double
G.COM.LG.F.64.X	Group compare less or greater floating-point double exact
G.COM.LG.F.128	Group compare less or greater floating-point quad
G.COM.LG.F.128.X	Group compare less or greater floating-point quad exact

Equivalencies

G.COM.G.F.16	Group compare greater floating-point half
G.COM.G.F.16.X	Group compare greater floating-point half exact
G.COM.G.F.32	Group compare greater floating-point single
G.COM.G.F.32.X	Group compare greater floating-point single exact
G.COM.G.F.64	Group compare greater floating-point double
G.COM.G.F.64.X	Group compare greater floating-point double exact
G.COM.G.F.128	Group compare greater floating-point quad
G.COM.G.F.128.X	Group compare greater floating-point quad exact
G.COM.LE.F.16	Group compare less equal floating-point half
G.COM.LE.F.16.X	Group compare less equal floating-point half exact
G.COM.LE.F.32	Group compare less equal floating-point single
G.COM.LE.F.32.X	Group compare less equal floating-point single exact
G.COM.LE.F.64	Group compare less equal floating-point double
G.COM.LE.F.64.X	Group compare less equal floating-point double exact
G.COM.LE.F.128	Group compare less equal floating-point quad
G.COM.LE.F.128.X	Group compare less equal floating-point quad exact

G.COM.G.F.prec rd,rc	→ G.COM.L.F.prec rc,rd
G.COM.G.F.prec.X rd,rc	→ G.COM.L.F.prec.X rc,rd
G.COM.LE.F.prec rd,rc	→ G.COM.GE.F.prec rc,rd
G.COM.LE.F.prec.X rd,rc	→ G.COM.GE.F.prec.X rc,rd


Selection

class	op	cond	type	prec	round/trap
set	COM	E LG L GE G LE	F	16 32 64 128	NONE X

Format

G.COM.op.prec.round rd,rc

rc=gcomopprecround(rd,rc)

Description

The contents of registers rd and rc are compared using the specified floating-point condition. If the result of the comparison is true for any corresponding pair of elements, a floating-point exception is raised. If a rounding option is specified, the operation raises a floating-point exception if a floating-point invalid operation occurs. If a rounding option is not specified, floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754.

Definition

```

def GroupCompareFloatingPoint(op,prec,round,rd,rc) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    for i ← 0 to 128-prec by prec
        di ← F(prec,di+prec-1..i)
        ci ← F(prec,ci+prec-1..i)
        if round≠NONE then
            if (di.t = SNAN) or (ci.t = SNAN) then
                raise FloatingPointArithmetic
            endif
            case op of
                G.COM.L.F, G.COM.G.E.F:
                    if (di.t = QNAN) or (ci.t = QNAN) then
                        raise FloatingPointArithmetic
                    endif
                others: //nothing
            endcase
        endif
        case op of
            G.COM.L.F:
                ai ← di?≥ci
            G.COM.G.E.F:
                ai ← di?<ci
            G.COM.E.F:
                ai ← di=ci
            G.COM.L.G.F:
                ai ← di≠ci
        endcase
        ai+prec-1..i ← ai
    endfor
    if (a ≠ 0) then
        raise FloatingPointArithmetic
    endif
enddef

```

Exceptions

Floating-point arithmetic

Group Copy Immediate

This operation copies an immediate value to a general register.

Operation codes

G.COPY.I.16	Group copy immediate doublet
G.COPY.I.32	Group signed copy immediate quadlet
G.COPY.I.64	Group signed copy immediate octlet
G.COPY.I.128	Group signed copy immediate hexlet

Equivalencies

G.COPY.I.8	Group copy immediate byte
G.SET	Group set
G.ZERO	Group zero

G.COPY.I.8 rd=(i ₃ i ₁ i _{7..0})	← G.COPY.I.16 rd=(0 i ₁ i _{7..0} i ₁ i _{7..0})
G.SET rd	← G.COPY.I.128 rd=-1
G.ZERO rd	← G.COPY.I.128 rd=0


Redundancies

G.COPY.I.size rd=-1	↔ G.SET rd
G.COPY.I.size rd=0	↔ G.ZERO rd

Format

G.COPY.I.size rd=i

rd=gcopy isize(i)



s ← i₁₆
imm ← i_{15..0}

Description

A 128-bit immediate value is produced from the operation code, the size field and the 16-bit imm field. The result is placed into register ra.

Definition

```
def GroupCopyImmediate(op,size,rd,imm) as
    s ← op[0]
```

```
case size of
 16:
  If s then
    ReservedInstruction
  endif
  a ← imm || imm
 32:
  a ← s16 || imm || s16 || imm || s16 || imm || s16 || imm
 64:
  a ← s48 || imm || s48 || imm
 128:
  a ← s112 || imm
endcase
RegWrite(rd, 128, a)
enddef
```

Exceptions

Reserved Instruction

Group Immediate

These operations take operands from a register and an immediate value, perform operations on partitions of bits in the operands, and place the concatenated results in a second register.

Operation codes

GADD.I.16	Group add immediate doublet
GADD.I.16.O	Group add immediate signed doublet check overflow
GADD.I.32	Group add immediate quadlet
GADD.I.32.O	Group add immediate signed quadlet check overflow
GADD.I.64	Group add immediate octlet
GADD.I.64.O	Group add immediate signed octlet check overflow
GADD.I.128	Group add immediate hexlet
GADD.I.128.O	Group add immediate signed hexlet check overflow
GADD.I.U.16.O	Group add immediate unsigned doublet check overflow
GADD.I.U.32.O	Group add immediate unsigned quadlet check overflow
GADD.I.U.64.O	Group add immediate unsigned octlet check overflow
GADD.I.U.128.O	Group add immediate unsigned hexlet check overflow
GAND.I.16	Group and immediate doublet
GAND.I.32	Group and immediate quadlet
GAND.I.64	Group and immediate octlet
GAND.I.128	Group and immediate hexlet
G.NAND.I.16	Group not and immediate doublet
G.NAND.I.32	Group not and immediate quadlet
G.NAND.I.64	Group not and immediate octlet
G.NAND.I.128	Group not and immediate hexlet
G.NOR.I.16	Group not or immediate doublet
G.NOR.I.32	Group not or immediate quadlet
G.NOR.I.64	Group not or immediate octlet
G.NOR.I.128	Group not or immediate hexlet
G.OR.I.16	Group or immediate doublet
G.OR.I.32	Group or immediate quadlet
G.OR.I.64	Group or immediate octlet
G.OR.I.128	Group or immediate hexlet
G.XOR.I.16	Group exclusive-or immediate doublet
G.XOR.I.32	Group exclusive-or immediate quadlet
G.XOR.I.64	Group exclusive-or immediate octlet
G.XOR.I.128	Group exclusive-or immediate hexlet

Equivalencies

G.ANDN.I.16	Group and not immediate doublet
G.ANDN.I.32	Group and not immediate quadlet
G.ANDN.I.64	Group and not immediate octlet
G.ANDN.I.128	Group and not immediate hexlet
G.COPY	Group copy
G.NOT	Group not
G.OR.N.I.16	Group or not immediate doublet
G.OR.N.I.32	Group or not immediate quadlet
G.OR.N.I.64	Group or not immediate octlet
G.OR.N.I.128	Group or not immediate hexlet
G.XNOR.I.16	Group exclusive-nor immediate doublet
G.XNOR.I.32	Group exclusive-nor immediate quadlet
G.XNOR.I.64	Group exclusive-nor immediate octlet
G.XNOR.I.128	Group exclusive-nor immediate hexlet

G.ANDN.I.size rd=rc,imm	→ G.AND.I.size rd=rc,-imm
G.COPY rd=rc	← G.OR.I.128 rd=rc,0
G.NOT rd=rc	← G.NOR.I.128 rd=rc,0
G.OR.N.I.size rd=rc,imm	→ G.OR.I.size rd=rc,-imm
G.XNOR.I.size rd=rc,imm	→ G.XOR.I.size rd=rc,-imm


Redundancies

G.ADD.I.size rd=rc,0	↔ G.COPY rd=rc
G.ADD.I.size.O rd=rc,0	↔ G.COPY rd=rc
G.ADD.I.U.size.O rd=rc,0	↔ G.COPY rd=rc
G.AND.I.size rd=rc,0	↔ G.ZERO rd
G.AND.I.size rd=rc,-1	↔ G.COPY rd=rc
G.NAND.I.size rd=rc,0	↔ G.SET rd
G.NAND.I.size rd=rc,-1	↔ G.NOT rd=rc
G.OR.I.size rd=rc,-1	↔ G.SET rd
G.NOR.I.size rd=rc,-1	↔ G.ZERO rd
G.XOR.I.size rd=rc,0	↔ G.COPY rd=rc
G.XOR.I.size rd=rc,-1	↔ G.NOT rd=rc

Format

op.size rd=rc,imm

rd=opsize(rc,imm)



$sz \leftarrow \log(size)-4$

Description

The contents of register rc is fetched, and a 128-bit immediate value is produced from the operation code, the size field and the 10-bit imm field. The specified operation is performed on these operands. The result is placed into register ra.

Definition

```
def GroupImmediate(op,size,rd,rc,imm) as
    c ← RegRead(rc, 128)
    s ← imm9
    case size of
        16:
            i16 ← s7 || imm
            b ← i16 || i16
        32:
            b ← s22 || imm || s22 || imm || s22 || imm || s22 || imm
        64:
            b ← s54 || imm || s54 || imm
        128:
            b ← s118 || imm
    endcase
    case op of
        G.AND.I:
            a ← c and b
        G.OR.I:
            a ← c or b
        G.NAND.I:
            a ← c nand b
        G.NOR.I:
            a ← c nor b
        G.XOR.I:
            a ← c xor b
        G.ADD.I:
            for i ← 0 to 128-size by size
                ai+size-1..i ← ci+size-1..i + bi+size-1..i
            endfor
        G.ADD.I.O:
            for i ← 0 to 128-size by size
                t ← (ci+size-1 || ci+size-1..i) + (bi+size-1 || bi+size-1..i)
                if tsize ≠ tsize-1 then
                    raise FixedPointArithmetic
                endif
                ai+size-1..i ← tsize-1..0
            endfor
        G.ADD.I.U.O:
            for i ← 0 to 128-size by size
                t ← (01 || ci+size-1..i) + (01 || bi+size-1..i)
                if tsize ≠ 0 then
                    raise FixedPointArithmetic
                endif
            endfor
```

```
     $a_{i+size-1..i} \leftarrow t_{size-1..0}$ 
  endfor
endcase
RegWrite(rd, 128, a)
enddef
```

Exceptions

Fixed-point arithmetic

Group Immediate Reversed

These operations take operands from a register and an immediate value, perform operations on partitions of bits in the operands, and place the concatenated results in a second register.

Operation codes

G.SET.AND.E.I.16	Group set and equal zero immediate doublets
G.SET.AND.E.I.32	Group set and equal zero immediate quadlets
G.SET.AND.E.I.64	Group set and equal zero immediate octlets
G.SET.AND.E.I.128	Group set and equal zero immediate hexlet
G.SET.AND.NE.I.16	Group set and not equal zero immediate doublets
G.SET.AND.NE.I.32	Group set and not equal zero immediate quadlets
G.SET.AND.NE.I.64	Group set and not equal zero immediate octlets
G.SET.AND.NE.I.128	Group set and not equal zero immediate hexlet
G.SET.E.I.16	Group set equal immediate doublets
G.SET.E.I.32	Group set equal immediate quadlets
G.SET.E.I.64	Group set equal immediate octlets
G.SET.E.I.128	Group set equal immediate hexlet
G.SET.GE.I.16	Group set greater equal immediate signed doublets
G.SET.GE.I.32	Group set greater equal immediate signed quadlets
G.SET.GE.I.64	Group set greater equal immediate signed octlets
G.SET.GE.I.128	Group set greater equal immediate signed hexlet
G.SET.GE.I.U.16	Group set greater equal immediate unsigned doublets
G.SET.GE.I.U.32	Group set greater equal immediate unsigned quadlets
G.SET.GE.I.U.64	Group set greater equal immediate unsigned octlets
G.SET.GE.I.U.128	Group set greater equal immediate unsigned hexlet
G.SET.LI.I.16	Group set signed less immediate doublets
G.SET.LI.I.32	Group set signed less immediate quadlets
G.SET.LI.I.64	Group set signed less immediate octlets
G.SET.LI.I.128	Group set signed less immediate hexlet
G.SET.LI.U.16	Group set less immediate signed doublets
G.SET.LI.U.32	Group set less immediate signed quadlets
G.SET.LI.U.64	Group set less immediate signed octlets
G.SET.LI.U.128	Group set less immediate signed hexlet
G.SET.NE.I.16	Group set not equal immediate doublets
G.SET.NE.I.32	Group set not equal immediate quadlets
G.SET.NE.I.64	Group set not equal immediate octlets
G.SET.NE.I.128	Group set not equal immediate hexlet
G.SUB.I.16	Group subtract immediate doublet
G.SUB.I.16.O	Group subtract immediate signed doublet check overflow
G.SUB.I.32	Group subtract immediate quadlet
G.SUB.I.32.O	Group subtract immediate signed quadlet check overflow
G.SUB.I.64	Group subtract immediate octlet
G.SUB.I.64.O	Group subtract immediate signed octlet check overflow
G.SUB.I.128	Group subtract immediat hexlet

G.SUB.I.128.O	Group subtract immediate signed hexlet check overflow
G.SUB.I.U.16.O	Group subtract immediate unsigned doublet check overflow
G.SUB.I.U.32.O	Group subtract immediate unsigned quadlet check overflow
G.SUB.I.U.64.O	Group subtract immediate unsigned octlet check overflow
G.SUB.I.U.128.O	Group subtract immediate unsigned hexlet check overflow

Equivalencies

G.NEG.16	Group negate doublet
G.NEG.16.O	Group negate signed doublet check overflow
G.NEG.32	Group negate quadlet
G.NEG.32.O	Group negate signed quadlet check overflow
G.NEG.64	Group negate octlet
G.NEG.64.O	Group negate signed octlet check overflow
G.NEG.128	Group negate hexlet
G.NEG.128.O	Group negate signed hexlet check overflow
G.SET.LE.I.16	Group set less equal immediate signed doublets
G.SET.LE.I.32	Group set less equal immediate signed quadlets
G.SET.LE.I.64	Group set less equal immediate signed octlets
G.SET.LE.I.128	Group set less equal immediate signed hexlet
G.SET.LE.I.U.16	Group set less equal immediate unsigned doublets
G.SET.LE.I.U.32	Group set less equal immediate unsigned quadlets
G.SET.LE.I.U.64	Group set less equal immediate unsigned octlets
G.SET.LE.I.U.128	Group set less equal immediate unsigned hexlet
G.SET.G.I.16	Group set immediate signed greater doublets
G.SET.G.I.32	Group set immediate signed greater quadlets
G.SET.G.I.64	Group set immediate signed greater octlets
G.SET.G.I.128	Group set immediate signed greater hexlet
G.SET.G.I.U.16	Group set greater immediate unsigned doublets
G.SET.G.I.U.32	Group set greater immediate unsigned quadlets
G.SET.G.I.U.64	Group set greater immediate unsigned octlets
G.SET.G.I.U.128	Group set greater immediate unsigned hexlet

G.NEG.size rd=rc	→ A.SUB.I.size rd=0,rc
G.NEG.size.O rd=rc	→ A.SUB.I.size.O rd=0,rc
G.SET.G.I.size rd=imm,rc	→ G.SET.GE.I.size rd=imm+1,rc
G.SET.G.I.U.size rd=imm,rc	→ G.SET.GE.I.U.size rd=imm+1,rc
G.SET.LE.I.size rd=imm,rc	→ G.SET.L.I.size rd=imm-1,rc
G.SET.LE.I.U.size rd=imm,rc	→ G.SET.L.I.U.size rd=imm-1,rc

Redundancies

G.SET.AND.E.I.size rd=rc,0	\Leftrightarrow	G.SET.size rd
G.SET.AND.NE.I.size rd=rc,0	\Leftrightarrow	G.ZERO rd
G.SET.AND.E.I.size rd=rc,-1	\Leftrightarrow	G.SET.E.Z.size rd=rc
G.SET.AND.NE.I.size rd=rc,-1	\Leftrightarrow	G.SET.NE.Z.size rd=rc
G.SET.E.I.size rd=rc,0	\Leftrightarrow	G.SET.E.Z.size rd=rc
G.SET.GE.I.size rd=rc,0	\Leftrightarrow	G.SET.GE.Z.size rd=rc
G.SET.L.I.size rd=rc,0	\Leftrightarrow	G.SET.L.Z.size rd=rc
G.SET.NE.I.size rd=rc,0	\Leftrightarrow	G.SET.NE.Z.size rd=rc
G.SET.GE.I.U.size rd=rc,0	\Leftrightarrow	G.SET.GE.U.Z.size rd=rc
G.SET.L.I.U.size rd=rc,0	\Leftrightarrow	G.SET.L.U.Z.size rd=rc


Selection

class	operation	cond	form	operand	size	check
arithmetic	SUB		I		16 32 64 128	
				NONE U	16 32 64 128	O
boolean	SET.AND	E	I		16 32 64 128	
	SET	NE				
	SET	L GE G LE	I	NONE U	16 32 64 128	

Format

op.size rd=imm,rc

rd=opsize{imm,rc}

sz $\leftarrow \log(\text{size}) - 4$ Description

The contents of register rc is fetched, and a 128-bit immediate value is produced from the operation code, the size field and the 10-bit imm field. The specified operation is performed on these operands. The result is placed into register rd.

Definition

```
def GroupImmediateReversed(op,size,ra,imm) as
    c ← RegRead(rc, 128)
    s ← imm9
    case size of
        16:
```

```

i16 ← s7 || imm
b ← i16 || i16
32:
b ← s22 || imm || s22 || imm || s22 || imm || s22 || imm
64:
b ← s54 || imm || s54 || imm
128:
b ← s118 || imm
endcase
case op of
  G.SUB.I:
    for i ← 0 to 128-size by size
      ai+size-1..i ← bi+size-1..i - ci+size-1..i
    endfor
  G.SUB.I.O:
    for i ← 0 to 128-size by size
      t ← (bi+size-1 || bi+size-1..i) - (ci+size-1 || ci+size-1..i)
      if (t.size ≠ t.size-1 then
          raise FixedPointArithmetic
      endif
      ai+size-1..i ← tsize-1..0
    endfor
  G.SUB.I.U.O:
    for i ← 0 to 128-size by size
      t ← (01 || bi+size-1..i) - (01 || ci+size-1..i)
      if (t.size ≠ 0 then
          raise FixedPointArithmetic
      endif
      ai+size-1..i ← tsize-1..0
    endfor
  G.SET.E.I:
    for i ← 0 to 128-size by size
      ai+size-1..i ← (bi+size-1..i = ci+size-1..i)size
    endfor
  G.SET.NE.I:
    for i ← 0 to 128-size by size
      ai+size-1..i ← (bi+size-1..i ≠ ci+size-1..i)size
    endfor
  G.SET.AND.E.I:
    for i ← 0 to 128-size by size
      ai+size-1..i ← ((bi+size-1..i and ci+size-1..i) = 0)size
    endfor
  G.SET.AND.NE.I:
    for i ← 0 to 128-size by size
      ai+size-1..i ← ((bi+size-1..i and ci+size-1..i) ≠ 0)size
    endfor
  G.SET.L.I:
    for i ← 0 to 128-size by size
      ai+size-1..i ← (bi+size-1..i < ci+size-1..i)size
    endfor
  G.SET.GE.I:
    for i ← 0 to 128-size by size

```

```
     $a_{i+size-1..i} \leftarrow (b_{i+size-1..i} \geq c_{i+size-1..i})^size$ 
  endfor
G.SET.LI.U:
  for i  $\leftarrow$  0 to 128-size by size
     $a_{i+size-1..i} \leftarrow ((0 \ 1 \ b_{i+size-1..i}) < (0 \ 1 \ c_{i+size-1..i}))^size$ 
  endfor
G.SET.GE.I.U:
  for i  $\leftarrow$  0 to 128-size by size
     $a_{i+size-1..i} \leftarrow ((0 \ 1 \ b_{i+size-1..i}) \geq (0 \ 1 \ c_{i+size-1..i}))^size$ 
  endfor
endcase
RegWrite(rd, 128, a)
enddef
```

Exceptions

Fixed-point arithmetic

Group Inplace

These operations take operands from three registers, perform operations on partitions of bits in the operands, and place the concatenated results in the third register.

Operation codes

GAAA.8	Group add add add bytes
GAAA.16	Group add add add doublets
GAAA.32	Group add add add quadlets
GAAA.64	Group add add add octlets
GAAA.128	Group add add add hexlet
GASA.8	Group add subtract add bytes
GASA.16	Group add subtract add doublets
GASA.32	Group add subtract add quadlets
GASA.64	Group add subtract add octlets
GASA.128	Group add subtract add hexlet

Equivalencies

GAAS.8	Group add add subtract bytes
GAAS.16	Group add add subtract doublets
GAAS.32	Group add add subtract quadlets
GAAS.64	Group add add subtract octlets
GAAS.128	Group add add subtract hexlet

GAAS.size rd@rc,rb	→ G.ASA.size rd@rb,rc
--------------------	-----------------------


Redundancies

GAAA.size rd@rc,rc	↔	G.SHLI.ADD.size rd=rd,rc,1
G.ASA.size rd@rc,rc	↔	G.NOP

Format

G.op.sizerd@rc,rb

rd=gopsizerd(rd,rc,rb)



Description

The contents of registers rd, rc and rb are fetched. The specified operation is performed on these operands. The result is placed into register rd.

Register rd is both a source and destination of this instruction.

Definition

```
def GroupInplace(op,size,rd,rc,rb) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-size by size
        case op of
            GAAA:
                 $a_{i+size-1..i} \leftarrow + d_{i+size-1..i} + c_{i+size-1..i} + b_{i+size-1..i}$ 
            GASA:
                 $a_{i+size-1..i} \leftarrow + d_{i+size-1..i} - c_{i+size-1..i} + b_{i+size-1..i}$ 
        endcase
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Group Reversed

These operations take two values from registers, perform operations on partitions of bits in the operands, and place the concatenated results in a register.

Operation codes

G.SET.AND.E.8	Group set and equal zero bytes
G.SET.AND.E.16	Group set and equal zero doublets
G.SET.AND.E.32	Group set and equal zero quadlets
G.SET.AND.E.64	Group set and equal zero octlets
G.SET.AND.E.128	Group set and equal zero hexlet
G.SET.AND.NE.8	Group set and not equal zero bytes
G.SET.AND.NE.16	Group set and not equal zero doublets
G.SET.AND.NE.32	Group set and not equal zero quadlets
G.SET.AND.NE.64	Group set and not equal zero octlets
G.SET.AND.NE.128	Group set and not equal zero hexlet
G.SET.E.8	Group set equal bytes
G.SET.E.16	Group set equal doublets
G.SET.E.32	Group set equal quadlets
G.SET.E.64	Group set equal octlets
G.SET.E.128	Group set equal hexlet
G.SET.GE.8	Group set greater equal signed bytes
G.SET.GE.16	Group set greater equal signed doublets
G.SET.GE.32	Group set greater equal signed quadlets
G.SET.GE.64	Group set greater equal signed octlets
G.SET.GE.128	Group set greater equal signed hexlet
G.SET.GE.U.8	Group set greater equal unsigned bytes
G.SET.GE.U.16	Group set greater equal unsigned doublets
G.SET.GE.U.32	Group set greater equal unsigned quadlets
G.SET.GE.U.64	Group set greater equal unsigned octlets
G.SET.GE.U.128	Group set greater equal unsigned hexlet
G.SET.L.8	Group set signed less bytes
G.SET.L.16	Group set signed less doublets
G.SET.L.32	Group set signed less quadlets
G.SET.L.64	Group set signed less octlets
G.SET.L.128	Group set signed less hexlet
G.SET.LU.8	Group set less unsigned bytes
G.SET.LU.16	Group set less unsigned doublets
G.SET.LU.32	Group set less unsigned quadlets
G.SET.LU.64	Group set less unsigned octlets
G.SET.LU.128	Group set less unsigned hexlet
G.SET.NE.8	Group set not equal bytes
G.SET.NE.16	Group set not equal doublets
G.SET.NE.32	Group set not equal quadlets
G.SET.NE.64	Group set not equal octlets

G.SET.NE.128	Group set not equal hexlet
G.SUB.8	Group subtract bytes
G.SUB.8.O	Group subtract signed bytes check overflow
G.SUB.16	Group subtract doublets
G.SUB.16.O	Group subtract signed doublets check overflow
G.SUB.32	Group subtract quadlets
G.SUB.32.O	Group subtract signed quadlets check overflow
G.SUB.64	Group subtract octlets
G.SUB.64.O	Group subtract signed octlets check overflow
G.SUB.128	Group subtract hexlet
G.SUB.128.O	Group subtract signed hexlet check overflow
G.SUB.L.8	Group subtract limit signed bytes
G.SUB.L.16	Group subtract limit signed doublets
G.SUB.L.32	Group subtract limit signed quadlets
G.SUB.L.64	Group subtract limit signed octlets
G.SUB.L.128	Group subtract limit signed hexlet
G.SUB.L.U.8	Group subtract limit unsigned bytes
G.SUB.L.U.16	Group subtract limit unsigned doublets
G.SUB.L.U.32	Group subtract limit unsigned quadlets
G.SUB.L.U.64	Group subtract limit unsigned octlets
G.SUB.L.U.128	Group subtract limit unsigned hexlet
G.SUB.U.8.O	Group subtract unsigned bytes check overflow
G.SUB.U.16.O	Group subtract unsigned doublets check overflow
G.SUB.U.32.O	Group subtract unsigned quadlets check overflow
G.SUB.U.64.O	Group subtract unsigned octlets check overflow
G.SUB.U.128.O	Group subtract unsigned hexlet check overflow

Equivalencies

<i>G.SET.E.Z.8</i>	Group set equal zero bytes
<i>G.SET.E.Z.16</i>	Group set equal zero doublets
<i>G.SET.E.Z.32</i>	Group set equal zero quadlets
<i>G.SET.E.Z.64</i>	Group set equal zero octlets
<i>G.SET.E.Z.128</i>	Group set equal zero hexlet
<i>G.SET.G.Z.8</i>	Group set greater zero signed bytes
<i>G.SET.G.Z.16</i>	Group set greater zero signed doublets
<i>G.SET.G.Z.32</i>	Group set greater zero signed quadlets
<i>G.SET.G.Z.64</i>	Group set greater zero signed octlets
<i>G.SET.G.Z.128</i>	Group set greater zero signed hexlet
<i>G.SET.GE.Z.8</i>	Group set greater equal zero signed bytes
<i>G.SET.GE.Z.16</i>	Group set greater equal zero signed doublets
<i>G.SET.GE.Z.32</i>	Group set greater equal zero signed quadlets
<i>G.SET.GE.Z.64</i>	Group set greater equal zero signed octlets
<i>G.SET.GE.Z.128</i>	Group set greater equal zero signed hexlet
<i>G.SET.L.Z.8</i>	Group set less zero signed bytes
<i>G.SET.L.Z.16</i>	Group set less zero signed doublets
<i>G.SET.L.Z.32</i>	Group set less zero signed quadlets
<i>G.SET.L.Z.64</i>	Group set less zero signed octlets
<i>G.SET.L.Z.128</i>	Group set less zero signed hexlet
<i>G.SET.LE.Z.8</i>	Group set less equal zero signed bytes
<i>G.SET.LE.Z.16</i>	Group set less equal zero signed doublets
<i>G.SET.LE.Z.32</i>	Group set less equal zero signed quadlets
<i>G.SET.LE.Z.64</i>	Group set less equal zero signed octlets
<i>G.SET.LE.Z.128</i>	Group set less equal zero signed hexlet
<i>G.SET.NE.Z.8</i>	Group set not equal zero bytes
<i>G.SET.NE.Z.16</i>	Group set not equal zero doublets
<i>G.SET.NE.Z.32</i>	Group set not equal zero quadlets
<i>G.SET.NE.Z.64</i>	Group set not equal zero octlets
<i>G.SET.NE.Z.128</i>	Group set not equal zero hexlet
<i>G.SET.LE.8</i>	Group set less equal signed bytes
<i>G.SET.LE.16</i>	Group set less equal signed doublets
<i>G.SET.LE.32</i>	Group set less equal signed quadlets
<i>G.SET.LE.64</i>	Group set less equal signed octlets
<i>G.SET.LE.128</i>	Group set less equal signed hexlet
<i>G.SET.LE.U.8</i>	Group set less equal unsigned bytes
<i>G.SET.LE.U.16</i>	Group set less equal unsigned doublets
<i>G.SET.LE.U.32</i>	Group set less equal unsigned quadlets
<i>G.SET.LE.U.64</i>	Group set less equal unsigned octlets
<i>G.SET.LE.U.128</i>	Group set less equal unsigned hexlet
<i>G.SET.G.8</i>	Group set signed greater bytes
<i>G.SET.G.16</i>	Group set signed greater doublets
<i>G.SET.G.32</i>	Group set signed greater quadlets
<i>G.SET.G.64</i>	Group set signed greater octlets

<i>G.SET.G.128</i>	Group set signed greater hexlet
<i>G.SET.G.U.8</i>	Group set greater unsigned bytes
<i>G.SET.G.U.16</i>	Group set greater unsigned doublets
<i>G.SET.G.U.32</i>	Group set greater unsigned quadlets
<i>G.SET.G.U.64</i>	Group set greater unsigned octlets
<i>G.SET.G.U.128</i>	Group set greater unsigned hexlet

<i>G.SET.E.Z.size rd=rc</i>	\leftarrow G.SET AND.E.size rd=rc,rc
<i>G.SET.G.Z.size rd=rc</i>	\Leftarrow G.SET.L.U.size rd=rc,rc
<i>G.SET.GE.Z.size rd=rc</i>	\Leftarrow G.SET.GE.size rd=rc,rc
<i>G.SET.L.Z.size rd=rc</i>	\Leftarrow G.SET.L.size rd=rc,rc
<i>G.SET.LE.Z.size rd=rc</i>	\Leftarrow G.SET.GE.U.size rd=rc,rc
<i>G.SET.NE.Z.size rd=rc</i>	\Leftarrow G.SET AND.NE.size rd=rc,rc
<i>G.SET.G.size rd=rb,rc</i>	\rightarrow G.SET.L.size rd=rc,rb
<i>G.SET.G.U.size rd=rb,rc</i>	\rightarrow G.SET.L.U.size rd=rc,rb
<i>G.SET.LE.size rd=rb,rc</i>	\rightarrow G.SET.GE.size rd=rc,rb
<i>G.SET.LE.U.size rd=rb,rc</i>	\rightarrow G.SET.GE.U.size rd=rc,rb

Redundancies

<i>G.SET.E.size rd=rc,rc</i>	\Leftrightarrow G.SET rd
<i>G.SET.NE.size rd=rc,rc</i>	\Leftrightarrow G.ZERO rd
<i>G.SUB.size rd=rc,rc</i>	\Leftrightarrow G.ZERO rd
<i>G.SUB.L.size rd=rc,rc</i>	\Leftrightarrow G.ZERO rd
<i>G.SUB.L.U.size rd=rc,rc</i>	\Leftrightarrow G.ZERO rd
<i>G.SUB.size.O rd=rc,rc</i>	\Leftrightarrow G.ZERO rd
<i>G.SUB.U.size.O rd=rc,rc</i>	\Leftrightarrow G.ZERO rd


Selection

class	operation	cond	operand	size	check
arithmetic	SUB			8 16 32 64 128	
			NONE U	8 16 32 64 128	O
	SUB.L		NONE U	8 16 32 64 128	
boolean	SET.AND	E		8 16 32 64 128	
	SET	NE		8 16 32 64 128	
	SET	L GE G LE	NONE U	8 16 32 64 128	
	SET	G GE L LE	Z	8 16 32 64 128	

Format

G.op.sizerd=rb,rc

rd=gopsizex{rb,rc}

Description

Two values are taken from the contents of registers rc and rb. The specified operation is performed, and the result is placed in register rd.

Definition

```

def GroupReversed(op,size,rd,rc,rb)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case op of
        G.SUB:
            for i ← 0 to 128-size by size
                ai+size-1..i ← bi+size-1..i - ci+size-1..i
            endfor
        G.SUB.L:
            for i ← 0 to 128-size by size
                t ← (bi+size-1 || bi+size-1..i) - (ci+size-1 || ci+size-1..i)
                ai+size-1..i ← (tsize ≠ tsize-1) ? (tsize || ti+size-1) : tsize-1..0
            endfor
        G.SUB.LU:
            for i ← 0 to 128-size by size
                t ← (01 || bi+size-1..i) - (01 || ci+size-1..i)
                ai+size-1..i ← (tsize ≠ 0) ? 0size: tsize-1..0
            endfor
        G.SUB.O:
            for i ← 0 to 128-size by size
                t ← (bi+size-1 || bi+size-1..i) - (ci+size-1 || ci+size-1..i)
                if (tsize ≠ tsize-1) then
                    raise FixedPointArithmetic
                endif
                ai+size-1..i ← tsize-1..0
            endfor
        G.SUB.U.O:
            for i ← 0 to 128-size by size
                t ← (01 || bi+size-1..i) - (01 || ci+size-1..i)
                if (tsize ≠ 0) then
                    raise FixedPointArithmetic
                endif
                ai+size-1..i ← tsize-1..0
            endfor
        G.SET.E:
    
```

```

        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow (b_{i+size-1..i} = c_{i+size-1..i})^{\text{size}}$ 
        endfor
G.SET.NE:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow (b_{i+size-1..i} \neq c_{i+size-1..i})^{\text{size}}$ 
        endfor
G.SET.AND.E:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow ((b_{i+size-1..i} \text{ and } c_{i+size-1..i}) = 0)^{\text{size}}$ 
        endfor
G.SET.AND.NE:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow ((b_{i+size-1..i} \text{ and } c_{i+size-1..i}) \neq 0)^{\text{size}}$ 
        endfor
G.SET.L:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow ((rc = rb) ? (b_{i+size-1..i} < 0) : (b_{i+size-1..i} < c_{i+size-1..i}))^{\text{size}}$ 
        endfor
G.SET.GE:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow ((rc = rb) ? (b_{i+size-1..i} \geq 0) : (b_{i+size-1..i} \geq c_{i+size-1..i}))^{\text{size}}$ 
        endfor
G.SET.L.U:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow ((rc = rb) ? (b_{i+size-1..i} > 0) :
                ((0 \mid\mid b_{i+size-1..i}) < (0 \mid\mid c_{i+size-1..i})))^{\text{size}}$ 
        endfor
G.SET.GE.U:
        for i ← 0 to 128-size by size
             $a_{i+size-1..i} \leftarrow ((rc = rb) ? (b_{i+size-1..i} \leq 0) :
                ((0 \mid\mid b_{i+size-1..i}) \geq (0 \mid\mid c_{i+size-1..i})))^{\text{size}}$ 
        endfor
    endcase
    RegWrite(rd, 128, a)
enddef

```

Exceptions

Fixed-point arithmetic

Group Reversed Floating-point

These operations take two values from registers, perform a group of floating-point arithmetic operations on partitions of bits in the operands, and place the concatenated results in a register.

Operation codes

G.SET.E.F.16	Group set equal floating-point half
G.SET.E.F.16.X	Group set equal floating-point half exact
G.SET.E.F.32	Group set equal floating-point single
G.SET.E.F.32.X	Group set equal floating-point single exact
G.SET.E.F.64	Group set equal floating-point double
G.SET.E.F.64.X	Group set equal floating-point double exact
G.SET.E.F.128	Group set equal floating-point quad
G.SET.E.F.128.X	Group set equal floating-point quad exact
G.SET.GE.F.16.X	Group set greater equal floating-point half exact
G.SET.GE.F.32.X	Group set greater equal floating-point single exact
G.SET.GE.F.64.X	Group set greater equal floating-point double exact
G.SET.GE.F.128.X	Group set greater equal floating-point quad exact
G.SET.LG.F.16	Group set less greater floating-point half
G.SET.LG.F.16.X	Group set less greater floating-point half exact
G.SET.LG.F.32	Group set less greater floating-point single
G.SET.LG.F.32.X	Group set less greater floating-point single exact
G.SET.LG.F.64	Group set less greater floating-point double
G.SET.LG.F.64.X	Group set less greater floating-point double exact
G.SET.LG.F.128	Group set less greater floating-point quad
G.SET.LG.F.128.X	Group set less greater floating-point quad exact
G.SET.L.F.16	Group set less floating-point half
G.SET.L.F.16.X	Group set less floating-point half exact
G.SET.L.F.32	Group set less floating-point single
G.SET.L.F.32.X	Group set less floating-point single exact
G.SET.L.F.64	Group set less floating-point double
G.SET.L.F.64.X	Group set less floating-point double exact
G.SET.L.F.128	Group set less floating-point quad
G.SET.L.F.128.X	Group set less floating-point quad exact
G.SET.GE.F.16	Group set greater equal floating-point half
G.SET.GE.F.32	Group set greater equal floating-point single
G.SET.GE.F.64	Group set greater equal floating-point double
G.SET.GE.F.128	Group set greater equal floating-point quad

Equivalencies

<i>G.SET.LE.F.16.X</i>	Group set less equal floating-point half exact
<i>G.SET.LE.F.32.X</i>	Group set less equal floating-point single exact
<i>G.SET.LE.F.64.X</i>	Group set less equal floating-point double exact
<i>G.SET.LE.F.128.X</i>	Group set less equal floating-point quad exact
<i>G.SET.G.F.16</i>	Group set greater floating-point half
<i>G.SET.G.F.16.X</i>	Group set greater floating-point half exact
<i>G.SET.G.F.32</i>	Group set greater floating-point single
<i>G.SET.G.F.32.X</i>	Group set greater floating-point single exact
<i>G.SET.G.F.64</i>	Group set greater floating-point double
<i>G.SET.G.F.64.X</i>	Group set greater floating-point double exact
<i>G.SET.G.F.128</i>	Group set greater floating-point quad
<i>G.SET.G.F.128.X</i>	Group set greater floating-point quad exact
<i>G.SET.LE.F.16</i>	Group set less equal floating-point half
<i>G.SET.LE.F.32</i>	Group set less equal floating-point single
<i>G.SET.LE.F.64</i>	Group set less equal floating-point double
<i>G.SET.LE.F.128</i>	Group set less equal floating-point quad

<i>G.SET.G.F.prec rd=rb,rc</i>	→ G.SET.L.F.prec rd=rc,rb
<i>G.SET.G.F.prec.X rd=rb,rc</i>	→ G.SET.L.F.prec.X rd=rc,rb
<i>G.SET.LE.F.prec rd=rb,rc</i>	→ G.SET.GE.F.prec rd=rc,rb
<i>G.SET.LE.F.prec.X rd=rb,rc</i>	→ G.SET.GE.F.prec.X rd=rc,rb


Selection

class	op	prec	round/trap
set	SET. E LG L GE G LE	16 32 64 128	NONE X

Format

G.op.prec.round rd=rb,rc

rc=gopprecround(rb,ra)

Description

The contents of registers ra and rb are combined using the specified floating-point operation. The result is placed in register rc. The operation is rounded using the specified rounding option or using round-to-nearest if not specified. If a rounding option is specified, the operation raises a floating-point exception if a floating-point invalid operation, divide by

zero, overflow, or underflow occurs, or when specified, if the result is inexact. If a rounding option is not specified, floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754.

Definition

```
def GroupFloatingPointReversed(op,prec,round,rd,rc,rb) as
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-prec by prec
        ci ← F[prec,ci+prec-1..i]
        bi ← F[prec,bi+prec-1..i]
        if round≠NONE then
            if (di.t = SNaN) or (ci.t = SNaN) then
                raise FloatingPointArithmetic
            endif
            case op of
                G.SET.L.F, G.SET.GE.F:
                    if (di.t = QNaN) or (ci.t = QNaN) then
                        raise FloatingPointArithmetic
                    endif
                others: //nothing
            endcase
        endif
        case op of
            G.SET.L.F:
                ai ← bi?≥ci
            G.SET.GE.F:
                ai ← bi?<ci
            G.SET.E.F:
                ai ← bi=ci
            G.SET.LG.F:
                ai ← bi≠ci
        endcase
        ai+prec-1..i ← ai
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

Floating-point arithmetic

Group Shift Left Immediate Add

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

G.SHL.I.ADD.8	Group shift left immediate add bytes
G.SHL.I.ADD.16	Group shift left immediate add doublets
G.SHL.I.ADD.32	Group shift left immediate add quadlets
G.SHL.I.ADD.64	Group shift left immediate add octlets
G.SHL.I.ADD.128	Group shift left immediate add hexlet


Redundancies

$$G.SHL.I.ADD.size\ rd=rd,rc,1 \Leftrightarrow GAAA.size\ rd@rc,rc$$

Format

G.op.sizerd=rc,rb,i

rd=gopsizerc(rc,rb,i)



assert $1 \leq i \leq 4$
 $sh \leftarrow i-1$

Description

The contents of registers rc and rb are partitioned into groups of operands of the size specified. Partitions of the contents of register rb are shifted left by the amount specified in the immediate field and added to partitions of the contents of register rc, yielding a group of results, each of which is the size specified. Overflows are ignored, and yield modular arithmetic results. The group of results is catenated and placed in register rd.

Definition

```
def GroupShiftLeftImmediateAdd(sh,size,ra,rb,rc)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-size by size
        ai+size-1..i ← ci+size-1..i + (bi+size-1-sh..i || 01+sh)
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Group Shift Left Immediate Subtract

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

G.SHL.I.SUB.8	Group shift left immediate subtract bytes
G.SHL.I.SUB.16	Group shift left immediate subtract doublets
G.SHL.I.SUB.32	Group shift left immediate subtract quadlets
G.SHL.I.SUB.64	Group shift left immediate subtract octlets
G.SHL.I.SUB.128	Group shift left immediate subtract hexlet


Redundancies

G.SHL.I.SUB.size rd=rc,1,rc	\Leftrightarrow	G.COPY rd=rc
-----------------------------	-------------------	--------------

Format

G.op.sizerd=rb,i,rc

rd=gopsiz(rb,i,rc)



assert $1 \leq i \leq 4$

sh \leftarrow i-1

Description

The contents of registers rc and rb are partitioned into groups of operands of the size specified. Partitions of the contents of register rc are subtracted from partitions of the contents of register rb shifted left by the amount specified in the immediate field, yielding a group of results, each of which is the size specified. Overflows are ignored, and yield modular arithmetic results. The group of results is catenated and placed in register rd.

Definition

```
def GroupShiftLeftImmediateSubtract(sh,size,ra,rb,rc)
    c  $\leftarrow$  RegRead(rc, 128)
    b  $\leftarrow$  RegRead(rb, 128)
    for i  $\leftarrow$  0 to 128-size by size
        ai+size-1..i  $\leftarrow$  (bi+size-1..i || 01+sh) - ci+size-1..i
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Group Subtract Halve

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

G.SUB.H.8.C	Group subtract halve signed bytes ceiling
G.SUB.H.8.F	Group subtract halve signed bytes floor
G.SUB.H.8.N	Group subtract halve signed bytes nearest
G.SUB.H.8.Z	Group subtract halve signed bytes zero
G.SUB.H.16.C	Group subtract halve signed doublets ceiling
G.SUB.H.16.F	Group subtract halve signed doublets floor
G.SUB.H.16.N	Group subtract halve signed doublets nearest
G.SUB.H.16.Z	Group subtract halve signed doublets zero
G.SUB.H.32.C	Group subtract halve signed quadlets ceiling
G.SUB.H.32.F	Group subtract halve signed quadlets floor
G.SUB.H.32.N	Group subtract halve signed quadlets nearest
G.SUB.H.32.Z	Group subtract halve signed quadlets zero
G.SUB.H.64.C	Group subtract halve signed octlets ceiling
G.SUB.H.64.F	Group subtract halve signed octlets floor
G.SUB.H.64.N	Group subtract halve signed octlets nearest
G.SUB.H.64.Z	Group subtract halve signed octlets zero
G.SUB.H.128.C	Group subtract halve signed hexlet ceiling
G.SUB.H.128.F	Group subtract halve signed hexlet floor
G.SUB.H.128.N	Group subtract halve signed hexlet nearest
G.SUB.H.128.Z	Group subtract halve signed hexlet zero
G.SUB.H.U.8.C	Group subtract halve unsigned bytes ceiling
G.SUB.H.U.8.F	Group subtract halve unsigned bytes floor
G.SUB.H.U.8.N	Group subtract halve unsigned bytes nearest
G.SUB.H.U.8.Z	Group subtract halve unsigned bytes zero
G.SUB.H.U.16.C	Group subtract halve unsigned doublets ceiling
G.SUB.H.U.16.F	Group subtract halve unsigned doublets floor
G.SUB.H.U.16.N	Group subtract halve unsigned doublets nearest
G.SUB.H.U.16.Z	Group subtract halve unsigned doublets zero
G.SUB.H.U.32.C	Group subtract halve unsigned quadlets ceiling
G.SUB.H.U.32.F	Group subtract halve unsigned quadlets floor
G.SUB.H.U.32.N	Group subtract halve unsigned quadlets nearest
G.SUB.H.U.32.Z	Group subtract halve unsigned quadlets zero
G.SUB.H.U.64.C	Group subtract halve unsigned octlets ceiling
G.SUB.H.U.64.F	Group subtract halve unsigned octlets floor
G.SUB.H.U.64.N	Group subtract halve unsigned octlets nearest
G.SUB.H.U.64.Z	Group subtract halve unsigned octlets zero
G.SUB.H.U.128.C	Group subtract halve unsigned hexlet ceiling
G.SUB.H.U.128.F	Group subtract halve unsigned hexlet floor
G.SUB.H.U.128.N	Group subtract halve unsigned hexlet nearest

G.SUB.H.U.128.Z	Group subtract halve unsigned hexlet zero
-----------------	---

Redundancies

G.SUB.H.size.rnd rd=rc,rc	\Leftrightarrow	G.ZERO rd
G.SUB.H.U.size.rnd rd=rc,rc	\Leftrightarrow	G.ZERO rd

Format

G.op.size.rnd rd=rb,rc

rd=gopsizernd(rb,rc)

31	24 23	18 17	12 11	6 5	21 0
	G.size	rd	rc	rb	op rnd

8 6 6 6 4 2

Description

The contents of registers rc and rb are partitioned into groups of operands of the size specified and subtracted, halved, rounded and limited as specified, yielding a group of results, each of which is the size specified. The group of results is catenated and placed in register rd.

The result of this operation is always signed, whether the operands are signed or unsigned.

Definition

```

def GroupSubtractHalve{op,rnd,size,rd,rc,rb}
  c ← RegRead[rc, 128]
  b ← RegRead[rb, 128]
  case op of
    G.SUB.H.C, G.SUB.H.F, G.SUB.H.N, G.SUB.H.Z:
      as ← cs ← bs ← 1
    G.SUB.H.U.C, G.SUB.H.U.F, G.SUB.H.U.N, G.SUB.H.U.Z
      as ← 1
      cs ← bs ← 0
  endcase
  for i ← 0 to 128-size by size
    p ← ((bs and bsize-1) || bsize-1+i..i) - ((cs and csize-1) || csize-1+i..i)
    case rnd of
      none, N:
        s ← 0size || ~p1
      Z:
        s ← 0size || psize
      F:
        s ← 0size+1
      C:
        s ← 0size || 11
    endcase
  endfor
enddef

```

```
v ← (as & psize) || p + (0 || s)
if vsize+1 = {as & vsize} then
    asize-1+i..i ← vsize..1
else
    asize-1+i..i ← as ? (vsize+1 || ~vsize-1) : 1size
endif
endfor
RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Group Ternary

These operations take three values from registers, perform a group of calculations on partitions of bits of the operands and place the catenated results in a fourth register.

Operation codes

G.MUX	Group multiplex
-------	-----------------


Redundancies

G.MUX ra=rd,rc,rc	\Leftrightarrow	G.COPY ra=rc
G.MUX ra=ra,rc,rb	\Leftrightarrow	G.BOOLEAN ra@rc,rb,0x11001010
G.MUX ra=rd,ra,rb	\Leftrightarrow	G.BOOLEAN ra@rd,rb,0x11100010
G.MUX ra=rd,rc,ra	\Leftrightarrow	G.BOOLEAN ra@rd,rc,0x11011000
G.MUX ra=rd,rd,rb	\Leftrightarrow	G.OR ra=rd,rb
G.MUX ra=rd,rc,rd	\Leftrightarrow	G.AND ra=rd,rc

Format

G.MUX ra=rd,rc,rb

ra=gmux(rd,rc,rb)



Description

The contents of registers rd, rc, and rb are fetched. Each bit of the result is equal to the corresponding bit of rc, if the corresponding bit of rd is set, otherwise it is the corresponding bit of rb. The result is placed into register ra.

Definition

```
def GroupTernary[op,size,rd,rc,rb,ra] as
  d ← RegRead(rd, 128)
  c ← RegRead(rc, 128)
  b ← RegRead(rb, 128)
  case op of
    G.MUX:
      a ← (c and d) or (b and not d)
    endcase
    RegWrite(ra, 128, a)
  enddef
```

Exceptions

none

Crossbar

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

X.COMPRESS.2	Crossbar compress signed pecks
X.COMPRESS.4	Crossbar compress signed nibbles
X.COMPRESS.8	Crossbar compress signed bytes
X.COMPRESS.16	Crossbar compress signed doublets
X.COMPRESS.32	Crossbar compress signed quadlets
X.COMPRESS.64	Crossbar compress signed octlets
X.COMPRESS.128	Crossbar compress signed hexlet
X.COMPRESS.U.2	Crossbar compress unsigned pecks
X.COMPRESS.U.4	Crossbar compress unsigned nibbles
X.COMPRESS.U.8	Crossbar compress unsigned bytes
X.COMPRESS.U.16	Crossbar compress unsigned doublets
X.COMPRESS.U.32	Crossbar compress unsigned quadlets
X.COMPRESS.U.64	Crossbar compress unsigned octlets
X.COMPRESS.U.128	Crossbar compress unsigned hexlet
X.EXPAND.2	Crossbar expand signed pecks
X.EXPAND.4	Crossbar expand signed nibbles
X.EXPAND.8	Crossbar expand signed bytes
X.EXPAND.16	Crossbar expand signed doublets
X.EXPAND.32	Crossbar expand signed quadlets
X.EXPAND.64	Crossbar expand signed octlets
X.EXPAND.128	Crossbar expand signed hexlet
X.EXPAND.U.2	Crossbar expand unsigned pecks
X.EXPAND.U.4	Crossbar expand unsigned nibbles
X.EXPAND.U.8	Crossbar expand unsigned bytes
X.EXPAND.U.16	Crossbar expand unsigned doublets
X.EXPAND.U.32	Crossbar expand unsigned quadlets
X.EXPAND.U.64	Crossbar expand unsigned octlets
X.EXPAND.U.128	Crossbar expand unsigned hexlet
X.ROTL.2	Crossbar rotate left pecks
X.ROTL.4	Crossbar rotate left nibbles
X.ROTL.8	Crossbar rotate left bytes
X.ROTL.16	Crossbar rotate left doublets
X.ROTL.32	Crossbar rotate left quadlets
X.ROTL.64	Crossbar rotate left octlets
X.ROTL.128	Crossbar rotate left hexlet
X.ROTR.2	Crossbar rotate right pecks
X.ROTR.4	Crossbar rotate right nibbles
X.ROTR.8	Crossbar rotate right bytes
X.ROTR.16	Crossbar rotate right doublets
X.ROTR.32	Crossbar rotate right quadlets

X.ROTR.64	Crossbar rotate right octlets
X.ROTR.128	Crossbar rotate right hexlet
X.SHL.2	Crossbar shift left pecks
X.SHL.2.O	Crossbar shift left signed pecks check overflow
X.SHL.4	Crossbar shift left nibbles
X.SHL.4.O	Crossbar shift left signed nibbles check overflow
X.SHL.8	Crossbar shift left bytes
X.SHL.8.O	Crossbar shift left signed bytes check overflow
X.SHL.16	Crossbar shift left doublets
X.SHL.16.O	Crossbar shift left signed doublets check overflow
X.SHL.32	Crossbar shift left quadlets
X.SHL.32.O	Crossbar shift left signed quadlets check overflow
X.SHL.64	Crossbar shift left octlets
X.SHL.64.O	Crossbar shift left signed octlets check overflow
X.SHL.128	Crossbar shift left hexlet
X.SHL.128.O	Crossbar shift left signed hexlet check overflow
X.SHL.U.2.O	Crossbar shift left unsigned pecks check overflow
X.SHL.U.4.O	Crossbar shift left unsigned nibbles check overflow
X.SHL.U.8.O	Crossbar shift left unsigned bytes check overflow
X.SHL.U.16.O	Crossbar shift left unsigned doublets check overflow
X.SHL.U.32.O	Crossbar shift left unsigned quadlets check overflow
X.SHL.U.64.O	Crossbar shift left unsigned octlets check overflow
X.SHL.U.128.O	Crossbar shift left unsigned hexlet check overflow
X.SHR.2	Crossbar signed shift right pecks
X.SHR.4	Crossbar signed shift right nibbles
X.SHR.8	Crossbar signed shift right bytes
X.SHR.16	Crossbar signed shift right doublets
X.SHR.32	Crossbar signed shift right quadlets
X.SHR.64	Crossbar signed shift right octlets
X.SHR.128	Crossbar signed shift right hexlet
X.SHR.U.2	Crossbar shift right unsigned pecks
X.SHR.U.4	Crossbar shift right unsigned nibbles
X.SHR.U.8	Crossbar shift right unsigned bytes
X.SHR.U.16	Crossbar shift right unsigned doublets
X.SHR.U.32	Crossbar shift right unsigned quadlets
X.SHR.U.64	Crossbar shift right unsigned octlets
X.SHR.U.128	Crossbar shift right unsigned hexlet

Selection

class	op	size
precision	EXPAND EXPAND.U COMPRESS COMPRESS.U	2 4 8 16 32 64 128
shift	ROTR ROTL SHR SHL SHL.O SHL.U.O SHR.U	2 4 8 16 32 64 128

Format

X.op.size rd=rc,rb

rd=xopsize(rc,rb)

31	252423	18 17	12 11	6 5	21 0	
XSHIFT	s	rd	rc	rb	op	sz
7	1	6	6	6	4	2

lsize $\leftarrow \log(\text{size})$ s $\leftarrow \text{lsize}_2$ sz $\leftarrow \text{lsize}_{1..0}$ Description

Two values are taken from the contents of registers rc and rb. The specified operation is performed, and the result is placed in register rd.

Definition

```
def Crossbar(op,size,rd,rc,rb)
    c  $\leftarrow$  RegRead[rc, 128]
    b  $\leftarrow$  RegRead[rb, 128]
    shift  $\leftarrow$  b and {size-1}
    case op5..2 11 02 of
        X.COMPRESS:
            hsize  $\leftarrow$  size/2
            for i  $\leftarrow$  0 to 64-hsize by hsize
                if shift  $\leq$  hsize then
                    ai+hsize-1..i  $\leftarrow$  ci+shift+hsize-1..i+shift
                else
                    ai+hsize-1..i  $\leftarrow$  ci+shift-hsizeshift-hsize || ci+hsize-1..i+shift
                endif
            endfor
            a127..64  $\leftarrow$  0
        X.COMPRESS.U:
            hsize  $\leftarrow$  size/2
            for i  $\leftarrow$  0 to 64-hsize by hsize
                if shift  $\leq$  hsize then
                    ai+hsize-1..i  $\leftarrow$  ci+shift+hsize-1..i+shift
                else
                    ai+hsize-1..i  $\leftarrow$  0shift-hsize || ci+hsize-1..i+shift
                endif
            endfor
            a127..64  $\leftarrow$  0
        X.EXPAND:
            hsize  $\leftarrow$  size/2
            for i  $\leftarrow$  0 to 64-hsize by hsize
                if shift  $\leq$  hsize then
                    ai+hsize-1..i+shift  $\leftarrow$  ci+hsize-1..ishift || ci+hsize-1..i+shift || 0shift
                endif
            endfor
            a127..64  $\leftarrow$  0
    endcase
enddef
```

```

        else
             $a_{i+size-1..i} \leftarrow c_{i+size-shift-1..i} \parallel 0^{shift}$ 
        endif
    endfor
X.EXPAND.U:
    hsize  $\leftarrow size/2$ 
    for i  $\leftarrow 0$  to  $64-hsize$  by hsize
        if shift  $\leq hsize$  then
             $a_{i+size-1..i} \leftarrow 0^{hsize-shift} \parallel c_{i+hsize-1..i} \parallel 0^{shift}$ 
        else
             $a_{i+size-1..i} \leftarrow c_{i+size-shift-1..i} \parallel 0^{shift}$ 
        endif
    endfor
X.ROTL:
    for i  $\leftarrow 0$  to  $128-size$  by size
         $a_{i+size-1..i} \leftarrow c_{i+size-1-shift..i} \parallel c_{i+size-1..i+size-1-shift}$ 
    endfor
X.ROTR:
    for i  $\leftarrow 0$  to  $128-size$  by size
         $a_{i+size-1..i} \leftarrow c_{i+shift-1..i} \parallel c_{i+size-1..i+shift}$ 
    endfor
X.SHL:
    for i  $\leftarrow 0$  to  $128-size$  by size
         $a_{i+size-1..i} \leftarrow c_{i+size-1-shift..i} \parallel 0^{shift}$ 
    endfor
X.SHL.O:
    for i  $\leftarrow 0$  to  $128-size$  by size
        if  $c_{i+size-1..i+size-1-shift} \neq c_{i+size-1-shift}^{shift+1}$  then
            raise FixedPointArithmetic
        endif
         $a_{i+size-1..i} \leftarrow c_{i+size-1-shift..i} \parallel 0^{shift}$ 
    endfor
X.SHL.U.O:
    for i  $\leftarrow 0$  to  $128-size$  by size
        if  $c_{i+size-1..i+size-shift} \neq 0^{shift}$  then
            raise FixedPointArithmetic
        endif
         $a_{i+size-1..i} \leftarrow c_{i+size-1-shift..i} \parallel 0^{shift}$ 
    endfor
X.SHR:
    for i  $\leftarrow 0$  to  $128-size$  by size
         $a_{i+size-1..i} \leftarrow c_{i+size-1}^{shift} \parallel c_{i+size-1..i+shift}$ 
    endfor
X.SHR.U:
    for i  $\leftarrow 0$  to  $128-size$  by size
         $a_{i+size-1..i} \leftarrow 0^{shift} \parallel c_{i+size-1..i+shift}$ 
    endfor
endcase
RegWrite(rd, 128, a)
enddef

```

Exceptions

Fixed-point arithmetic

Crossbar Extract

These operations take operands from three registers, perform operations on partitions of bits in the operands, and place the concatenated results in a fourth register.


Operation codes

X.EXTRACT	Crossbar extract
-----------	------------------

Format

X.EXTRACT ra=rd,rc,rb


ra=xextract(rd,rc,rb)



Description

The contents of registers rd, rc, and rb are fetched. The specified operation is performed on these operands. The result is placed into register ra.

Bits 31..0 of the contents of register rb specifies several parameters which control the manner in which data is extracted, and for certain operations, the manner in which the operation is performed. The position of the control fields allows for the source position to be added to a fixed control value for dynamic computation, and allows for the lower 16 bits of the control field to be set for some of the simpler extract cases by a single GCOPYI.128 instruction. The control fields are further arranged so that if only the low order 8 bits are non-zero, a 128-bit extraction with truncation and no rounding is performed.



The table below describes the meaning of each label:


label	bits	meaning
fsize	8	field size
dpos	8	destination position
x	1	reserved
s	1	signed vs. unsigned
n	1	reserved
m	1	merge vs. extract
l	1	reserved
rnd	2	reserved
gssp	9	group size and source position

The 9-bit **gssp** field encodes both the group size, **gsize**, and source position, **spos**, according to the formula $\text{gssp} = 512 \cdot 4^{\text{gsize}} + \text{spos}$. The group size, **gsize**, is a power of two in the range 1..128. The source position, **spos**, is in the range 0.. $(2^{\text{gsize}})-1$.

The values in the **s**, **n**, **m**, **l**, and **rnd** fields have the following meaning:


values	s	n	m	l	rnd
0	unsigned				
1	signed				
2			extract		
3			merge		

For the X.EXTRACT instruction, when **m=0**, the parameters are interpreted to select **fields** from the catenated contents of registers **rd** and **rc**, extracting values which are catenated and placed in register **ra**:



Crossbar extract

For a crossbar-merge-extract (X.EXTRACT when **m=1**), the parameters are interpreted to merge fields from the contents of register **rd** with the contents of register **rc**. The results are catenated and placed in register **ra**.



Crossbar merge extract

Definition

```

def CrossbarExtract(op,ra,rb,rc,rd) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case b8..0 of
        0..255:
            gsize ← 128
        256..383:
            gsize ← 64
        384..447:
            gsize ← 32
        448..479:
            gsize ← 16
        480..495:
            gsize ← 8
        496..503:
            gsize ← 4
        504..507:
            gsize ← 2
        508..511:
            gsize ← 1
    endcase
    m ← b12
    as ← signed ← b14
    h ← (2-m)*gsize
    spos ← (b8..0) and ((2-m)*gsize-1)
    dpos ← (0 || b23..16) and (gsize-1)
    sfsz ← (0 || b31..24) and (gsize-1)
    tfsz ← (sfsz = 0) or ((sfsz+dpos) > gsize) ? gsize-dpos : sfsz
    fsize ← (tfsz + spos > h) ? h - spos : tfsz
    for i ← 0 to 128-gsize by gsize
        case op of
            X.EXTRACT:
                if m then
                    p ← dgsz+i..i
                else
                    p ← (d || c)2*(gsz+i)..2*i
                endif
        endcase
        v ← (as & ph-1) || p
        w ← (as & vspos+fsize-1)gsize-fsize-dpos || vfsize-1+spos..spos || 0dpos
        if m then
            asize-1+i..i ← cgsz-1+i..dpos+fsize+i || wdpos+fsize-1..dpos || cdpos-1+1..i
        else
            asize-1+i..i ← w
        endif
    endfor
    RegWrite(ra, 128, a)
enddef

```

Exceptions

none

Crossbar Field

These operations take operands from a register and two immediate values, perform operations on partitions of bits in the operands, and place the concatenated results in the second register.

Operation codes

X.DEPPOSIT.2	Crossbar deposit signed pecks
X.DEPPOSIT.4	Crossbar deposit signed nibbles
X.DEPPOSIT.8	Crossbar deposit signed bytes
X.DEPPOSIT.16	Crossbar deposit signed doublets
X.DEPPOSIT.32	Crossbar deposit signed quadlets
X.DEPPOSIT.64	Crossbar deposit signed octlets
X.DEPPOSIT.128	Crossbar deposit signed hexlet
X.DEPPOSIT.U.2	Crossbar deposit unsigned pecks
X.DEPPOSIT.U.4	Crossbar deposit unsigned nibbles
X.DEPPOSIT.U.8	Crossbar deposit unsigned bytes
X.DEPPOSIT.U.16	Crossbar deposit unsigned doublets
X.DEPPOSIT.U.32	Crossbar deposit unsigned quadlets
X.DEPPOSIT.U.64	Crossbar deposit unsigned octlets
X.DEPPOSIT.U.128	Crossbar deposit unsigned hexlet
X.WITHDRAW.U.2	Crossbar withdraw unsigned pecks
X.WITHDRAW.U.4	Crossbar withdraw unsigned nibbles
X.WITHDRAW.U.8	Crossbar withdraw unsigned bytes
X.WITHDRAW.U.16	Crossbar withdraw unsigned doublets
X.WITHDRAW.U.32	Crossbar withdraw unsigned quadlets
X.WITHDRAW.U.64	Crossbar withdraw unsigned octlets
X.WITHDRAW.U.128	Crossbar withdraw unsigned hexlet
X.WITHDRAW.2	Crossbar withdraw pecks
X.WITHDRAW.4	Crossbar withdraw nibbles
X.WITHDRAW.8	Crossbar withdraw bytes
X.WITHDRAW.16	Crossbar withdraw doublets
X.WITHDRAW.32	Crossbar withdraw quadlets
X.WITHDRAW.64	Crossbar withdraw octlets
X.WITHDRAW.128	Crossbar withdraw hexlet

Equivalencies

X.SEX.I.2	Crossbar extend immediate signed pecks
X.SEX.I.4	Crossbar extend immediate signed nibbles
X.SEX.I.8	Crossbar extend immediate signed bytes
X.SEX.I.16	Crossbar extend immediate signed doublets
X.SEX.I.32	Crossbar extend immediate signed quadlets
X.SEX.I.64	Crossbar extend immediate signed octlets
X.SEX.I.128	Crossbar extend immediate signed hexlet
X.ZEX.I.2	Crossbar extend immediate unsigned pecks
X.ZEX.I.4	Crossbar extend immediate unsigned nibbles
X.ZEX.I.8	Crossbar extend immediate unsigned bytes
X.ZEX.I.16	Crossbar extend immediate unsigned doublets
X.ZEX.I.32	Crossbar extend immediate unsigned quadlets
X.ZEX.I.64	Crossbar extend immediate unsigned octlets
X.ZEX.I.128	Crossbar extend immediate unsigned hexlet

X.SHL.I.gsize rd=rc,i	→ X.DEPOSIT.gsize rd=rc,size-i,i
X.SHR.I.gsize rd=rc,i	→ X.WITHDRAW.gsize rd=rc,size-i,i
X.SHRU.I.gsize rd=rc,i	→ X.WITHDRAW.U.gsize rd=rc,size-i,i
X.SEX.I.gsize rd=rc,i	→ X.DEPOSIT.gsize rd=rc,i,0
X.ZEX.I.gsize rd=rc,i	→ X.DEPOSIT.U.gsize rd=rc,i,0


Redundancies

X.DEPOSIT.gsize rd=rc,gsize,0	↔ X.COPY rd=rc
X.DEPOSIT.U.gsize rd=rc,gsize,0	↔ X.COPY rd=rc
X.WITHDRAW.gsize rd=rc,gsize,0	↔ X.COPY rd=rc
X.WITHDRAW.U.gsize rd=rc,gsize,0	↔ X.COPY rd=rc

Format

X.op.gsize rd=rc,isize,ishift

rd=xopgsize(rc,isize,ishift)



```


assert isize+ishift ≤ gsize
assert isize≥1
ih0 || gsfs ← 128-gsize+isize-1
ih1 || gsfp ← 128-gsize+ishift

```

Description


The contents of register rc is fetched, and 7-bit immediate values are taken from the 2-bit ih and the 6-bit gsfp and gfsf fields. The specified operation is performed on these operands. The result is placed into register rd.

The diagram below shows legal values for the ih, gsfp and gfsf fields, indicating the group size to which they apply.




The ih, gsfp and gfsf fields encode three values: the group size, the field size, and a shift amount. The shift amount can also be considered to be the source bit field position for group-withdraw instructions or the destination bit field position for group-deposit instructions. The encoding is designed so that combining the gsfp and gfsf fields with a bitwise-and produces a result which can be decoded to the group size, and so the field size and shift amount can be easily decoded once the group size has been determined.

The crossbar-deposit instructions deposit a bit field from the lower bits of each group partition of the source to a specified bit position in the result. The value is either sign-extended or zero-extended, as specified.



The crossbar-withdraw instructions withdraw a bit field from a specified bit position in the each group partition of the source and place it in the lower bits in the result. The value is either sign-extended or zero-extended, as specified.



Definition

```
def CrossbarField(op,rd,rc,gsfp,gsfs) as
  c ← RegRead[rc, 128]
  case [(op1 || gsfp) and (op0 || gsfs)] of
    0..63:
      gsize ← 128
    64..95:
      gsize ← 64
```

```

96..111:
    gsize ← 32
112..119:
    gsize ← 16
120..123:
    gsize ← 8
124..125:
    gsize ← 4
126:
    gsize ← 2
127:
    raise ReservedInstruction
endcase
ishift ← (op1 || gsfp) and (gsize-1)
 isize ← ((op0 || gsfp) and (gsize-1))+1
if (ishift>isize)
    raise ReservedInstruction
endif
case op of
    X.DEPOSIT:
        for i ← 0 to 128-gsize by gsize
            ai+gsize-1..i ← ci+gsize-isize-1..i+gsize-1..i || 0ishift
        endfor
    X.DEPOSIT.U:
        for i ← 0 to 128-gsize by gsize
            ai+gsize-1..i ← 0gsize-isize-1..i || ci+gsize-1..i || 0ishift
        endfor
    X.WITHDRAW:
        for i ← 0 to 128-gsize by gsize
            ai+gsize-1..i ← ci+gsize+ishift-1..i+gsize+ishift-1..i+ishift
        endfor
    X.WITHDRAW.U:
        for i ← 0 to 128-gsize by gsize
            ai+gsize-1..i ← 0gsize-isize || ci+gsize+ishift-1..i+ishift
        endfor
    endcase
    RegWrite(rd, 128, a)
enddef

```

Exceptions

Reserved instruction

Crossbar Field Inplace

These operations take operands from two registers and two immediate values, perform operations on partitions of bits in the operands, and place the concatenated results in the second register.

Operation codes

X.DEPPOSIT.M.2	Crossbar deposit merge pecks
X.DEPPOSIT.M.4	Crossbar deposit merge nibbles
X.DEPPOSIT.M.8	Crossbar deposit merge bytes
X.DEPPOSIT.M.16	Crossbar deposit merge doublets
X.DEPPOSIT.M.32	Crossbar deposit merge quadlets
X.DEPPOSIT.M.64	Crossbar deposit merge octlets
X.DEPPOSIT.M.128	Crossbar deposit merge hexlet

Equivalencies

X.DEPPOSIT.M.1	Crossbar deposit merge bits
X.DEPPOSIT.M.1 rd@rc,1,0	→ X.COPY rd=rc


Redundancies

X.DEPPOSIT.M.gsize rd@rc,gsize,0	↔	X.COPY rd=rc
----------------------------------	---	--------------

Format

X.op.gsize rd@rc, isize, ishift

rd=xopgsize(rd,rc,isize,ishift)



```


assert isize+ishift ≤ gsize
assert isize≥1
ih0 11 gsfs ← 128-gsize+isize-1
ih1 11 gsfp ← 128-gsize+ishift

```

Description


The contents of registers rd and rc are fetched, and 7-bit immediate values are taken from the 2-bit lh and the 6-bit gsfp and gsfs fields. The specified operation is performed on these operands. The result is placed into register rd.

The diagram below shows legal values for the ih, gsfp and gfsf fields, indicating the group size to which they apply.



The ih, gsfp and gfsf fields encode three values: the group size, the field size, and a shift amount. The shift amount can also be considered to be the source bit field position for group-withdraw instructions or the destination bit field position for group-deposit instructions. The encoding is designed so that combining the gsfp and gfsf fields with a bitwise-and produces a result which can be decoded to the group size, and so the field size and shift amount can be easily decoded once the group size has been determined.

The crossbar-deposit-merge instructions deposit a bit field from the lower bits of each group partition of the source to a specified bit position in the result. The value is merged with the contents of register rd at bit positions above and below the deposited bit field. No sign- or zero-extension is performed by this instruction.



Definition

```

def CrossbarFieldInplace(op,rd,rc,gsfp,gsfs) as
  c ← RegRead(rc, 128)
  d ← RegRead(rd, 128)
  case {op1 || gsfp} and {op0 || gsfs} of
    0..63:
      gsize ← 128
    64..95:
      gsize ← 64
    96..111:
      gsize ← 32
    112..119:
      gsize ← 16
    120..123:
      gsize ← 8
    124..125:
      gsize ← 4
    126:
      gsize ← 2
    127:
      raise ReservedInstruction
  endcase
  ishift ← {op1 || gsfp} and {gsize-1}
  isize ← {op0 || gsfs} and {gsize-1}+1
  if (ishift+isize>gsize)
    raise ReservedInstruction
  endif
  for i ← 0 to 128-gsize by gsize
    a[i:gsize-1..i] ← d[i:gsize-1..i+isize+ishift] || c[i:gsize-1..i] || d[i:ishift-1..i]
  endfor
  RegWrite(rd, 128, a)
enddef

```

Exceptions

Reserved instruction

Crossbar Inplace

These operations take operands from three registers, perform operations on partitions of bits in the operands, and place the concatenated results in the third register.


Operation codes

X.SHL.M.2	Crossbar shift left merge pecks
X.SHL.M.4	Crossbar shift left merge nibbles
X.SHL.M.8	Crossbar shift left merge bytes
X.SHL.M.16	Crossbar shift left merge doublets
X.SHL.M.32	Crossbar shift left merge quadlets
X.SHL.M.64	Crossbar shift left merge octlets
X.SHL.M.128	Crossbar shift left merge hexlet
X.SHR.M.2	Crossbar shift right merge pecks
X.SHR.M.4	Crossbar shift right merge nibbles
X.SHR.M.8	Crossbar shift right merge bytes
X.SHR.M.16	Crossbar shift right merge doublets
X.SHR.M.32	Crossbar shift right merge quadlets
X.SHR.M.64	Crossbar shift right merge octlets
X.SHR.M.128	Crossbar shift right merge hexlet

Format

X.op.size rd@rc,rb

rd=xopsize(rd,rc,rb)



lsize $\leftarrow \log(\text{size})$

s $\leftarrow \text{lsize}_2$

sz $\leftarrow \text{lsize}_{1..0}$

Description

The contents of registers rd, rc and rb are fetched. The specified operation is performed on these operands. The result is placed into register rd.

Register rd is both a source and destination of this instruction.

Definition

```
def CrossbarInplace{op,size,rd,rc,rb} as
  d  $\leftarrow \text{RegRead}(rd, 128)$ 
  c  $\leftarrow \text{RegRead}(rc, 128)$ 
```

```
b ← RegRead(rb, 128)
shift ← b and (size-1)
for i ← 0 to 128-size by size
    case op of
        X.SHR.M:
            ai+size-1..i ← ci+shift-1..i || di+size-1..i+shift
        X.SHL.M:
            ai+size-1..i ← di+size-1-shift..i || ci+shift-1..i
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Crossbar Short Immediate

These operations take operands from a register and a short immediate value, perform operations on partitions of bits in the operands, and place the concatenated results in a register.

Operation codes

X.COMPRESS.I.2	Crossbar compress immediate signed pecks
X.COMPRESS.I.4	Crossbar compress immediate signed nibbles
X.COMPRESS.I.8	Crossbar compress immediate signed bytes
X.COMPRESS.I.16	Crossbar compress immediate signed doublets
X.COMPRESS.I.32	Crossbar compress immediate signed quadlets
X.COMPRESS.I.64	Crossbar compress immediate signed octlets
X.COMPRESS.I.128	Crossbar compress immediate signed hexlet
X.COMPRESS.I.U.2	Crossbar compress immediate unsigned pecks
X.COMPRESS.I.U.4	Crossbar compress immediate unsigned nibbles
X.COMPRESS.I.U.8	Crossbar compress immediate unsigned bytes
X.COMPRESS.I.U.16	Crossbar compress immediate unsigned doublets
X.COMPRESS.I.U.32	Crossbar compress immediate unsigned quadlets
X.COMPRESS.I.U.64	Crossbar compress immediate unsigned octlets
X.COMPRESS.I.U.128	Crossbar compress immediate unsigned hexlet
X.EXPAND.I.2	Crossbar expand immediate signed pecks
X.EXPAND.I.4	Crossbar expand immediate signed nibbles
X.EXPAND.I.8	Crossbar expand immediate signed bytes
X.EXPAND.I.16	Crossbar expand immediate signed doublets
X.EXPAND.I.32	Crossbar expand immediate signed quadlets
X.EXPAND.I.64	Crossbar expand immediate signed octlets
X.EXPAND.I.128	Crossbar expand immediate signed hexlet
X.EXPAND.I.U.2	Crossbar expand immediate unsigned pecks
X.EXPAND.I.U.4	Crossbar expand immediate unsigned nibbles
X.EXPAND.I.U.8	Crossbar expand immediate unsigned bytes
X.EXPAND.I.U.16	Crossbar expand immediate unsigned doublets
X.EXPAND.I.U.32	Crossbar expand immediate unsigned quadlets
X.EXPAND.I.U.64	Crossbar expand immediate unsigned octlets
X.EXPAND.I.U.128	Crossbar expand immediate unsigned hexlet
X.ROTL.I.2	Crossbar rotate left immediate pecks
X.ROTL.I.4	Crossbar rotate left immediate nibbles
X.ROTL.I.8	Crossbar rotate left immediate bytes
X.ROTL.I.16	Crossbar rotate left immediate doublets
X.ROTL.I.32	Crossbar rotate left immediate quadlets
X.ROTL.I.64	Crossbar rotate left immediate octlets
X.ROTL.I.128	Crossbar rotate left immediate hexlet
X.ROTR.I.2	Crossbar rotate right immediate pecks
X.ROTR.I.4	Crossbar rotate right immediate nibbles
X.ROTR.I.8	Crossbar rotate right immediate bytes
X.ROTR.I.16	Crossbar rotate right immediate doublets

X.ROTR.I.32	Crossbar rotate right immediate quadlets
X.ROTR.I.64	Crossbar rotate right immediate octlets
X.ROTR.I.128	Crossbar rotate right immediate hexlet
X.SHL.I.2	Crossbar shift left immediate pecks
X.SHL.I.2.O	Crossbar shift left immediate signed pecks check overflow
X.SHL.I.4	Crossbar shift left immediate nibbles
X.SHL.I.4.O	Crossbar shift left immediate signed nibbles check overflow
X.SHL.I.8	Crossbar shift left immediate bytes
X.SHL.I.8.O	Crossbar shift left immediate signed bytes check overflow
X.SHL.I.16	Crossbar shift left immediate doublets
X.SHL.I.16.O	Crossbar shift left immediate signed doublets check overflow
X.SHL.I.32	Crossbar shift left immediate quadlets
X.SHL.I.32.O	Crossbar shift left immediate signed quadlets check overflow
X.SHL.I.64	Crossbar shift left immediate octlets
X.SHL.I.64.O	Crossbar shift left immediate signed octlets check overflow
X.SHL.I.128	Crossbar shift left immediate hexlet
X.SHL.I.128.O	Crossbar shift left immediate signed hexlet check overflow
X.SHL.I.U.2.O	Crossbar shift left immediate unsigned pecks check overflow
X.SHL.I.U.4.O	Crossbar shift left immediate unsigned nibbles check overflow
X.SHL.I.U.8.O	Crossbar shift left immediate unsigned bytes check overflow
X.SHL.I.U.16.O	Crossbar shift left immediate unsigned doublets check overflow
X.SHL.I.U.32.O	Crossbar shift left immediate unsigned quadlets check overflow
X.SHL.I.U.64.O	Crossbar shift left immediate unsigned octlets check overflow
X.SHL.I.U.128.O	Crossbar shift left immediate unsigned hexlet check overflow
X.SHR.I.2	Crossbar signed shift right immediate pecks
X.SHR.I.4	Crossbar signed shift right immediate nibbles
X.SHR.I.8	Crossbar signed shift right immediate bytes
X.SHR.I.16	Crossbar signed shift right immediate doublets
X.SHR.I.32	Crossbar signed shift right immediate quadlets
X.SHR.I.64	Crossbar signed shift right immediate octlets
X.SHR.I.128	Crossbar signed shift right immediate hexlet
X.SHR.I.U.2	Crossbar shift right immediate unsigned pecks
X.SHR.I.U.4	Crossbar shift right immediate unsigned nibbles
X.SHR.I.U.8	Crossbar shift right immediate unsigned bytes
X.SHR.I.U.16	Crossbar shift right immediate unsigned doublets
X.SHR.I.U.32	Crossbar shift right immediate unsigned quadlets
X.SHR.I.U.64	Crossbar shift right immediate unsigned octlets
X.SHR.I.U.128	Crossbar shift right immediate unsigned hexlet

Equivalencies

X.COPY	Crossbar copy
X.NOP	Crossbar no operation
X.COPY rd=rc	← X.ROTL.I.128 rd=rc,0
X.NOP	← X.COPY r0=r0

Redundancies

$X.ROTL.I.gsize\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$
$X.ROTR.I.gsize\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$
$X.ROTR.I.gsize\ rd=rc,shift$	$\Leftrightarrow X.ROTL.I.gsize\ rd=rc,gsize-shift$
$X.SHL.I.gsize\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$
$X.SHL.I.gsize.O\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$
$X.SHL.I.U.gsize.O\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$
$X.SHR.I.gsize\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$
$X.SHR.I.U.gsize\ rd=rc,0$	$\Leftrightarrow X.COPY\ rd=rc$


Selection

class	op	size
precision	COMPRESS.I COMPRESS.I.U EXPAND.I EXPAND.I.U	2 4 8 16 32 64 128
shift	ROTL.I ROTR.I SHL.I SHL.I.O SHL.I.U.O SHR.I SHR.I.U	2 4 8 16 32 64 128
copy	COPY	

Format

X.op.size rd=rc,shift

rd=xopsizex(rc,shift)

 $t \leftarrow 256 - 2 * \text{size} + \text{shift}$ $\text{op}_{1..0} \leftarrow t_{7..6}$ $\text{simm} \leftarrow t_{5..0}$ Description

A 128-bit value is taken from the contents of register rc. The second operand is taken from simm. The specified operation is performed, and the result is placed in register rd.

Definition

```
def CrossbarShortImmediate(op,rd,rc,simm)
    case {op1..0 11 simm} of
        0..127:
            size ← 128
        128..191:
            size ← 64
```

```

192..223:
    size ← 32
224..239:
    size ← 16
240..247:
    size ← 8
248..251:
    size ← 4
252..253:
    size ← 2
254..255:
    raise ReservedInstruction
endcase
shift ← (op0 || simm) and (size-1)
c ← RegRead(rc, 128)
case (op5..2 || 02) of
    X.COMPRESS.I:
        hsize ← size/2
        for i ← 0 to 64-hsize by hsize
            if shift ≤ hsize then
                ai+hsize-1..i ← ci+shift+hsize-1..i+shift
            else
                ai+hsize-1..i ← ci+shift+hsize-1..i+shift || ci+size-1..i+shift
            endif
        endfor
        a127..64 ← 0
    X.COMPRESS.I.U:
        hsize ← size/2
        for i ← 0 to 64-hsize by hsize
            if shift ≤ hsize then
                ai+hsize-1..i ← ci+shift+hsize-1..i+shift
            else
                ai+hsize-1..i ← 0shift-hsize || ci+size-1..i+shift
            endif
        endfor
        a127..64 ← 0
    X.EXPAND.I:
        hsize ← size/2
        for i ← 0 to 64-hsize by hsize
            if shift ≤ hsize then
                ai+hsize-1..i ← ci+shift+hsize-1..i || 0shift
            else
                ai+hsize-1..i ← ci+size-shift-1..i || 0shift
            endif
        endfor
    X.EXPAND.I.U:
        hsize ← size/2
        for i ← 0 to 64-hsize by hsize
            if shift ≤ hsize then
                ai+hsize-1..i ← 0shift-hsize || ci+hsize-1..i || 0shift
            else
                ai+hsize-1..i ← ci+size-shift-1..i || 0shift
            endif
        endfor

```

```

        endif
    endfor
X.SHL.I:
    for i ← 0 to 128-size by size
         $a_{i+size-1..i} \leftarrow c_{i+size-1..i+shift..i} \parallel 0^{shift}$ 
    endfor
X.SHL.I.O:
    for i ← 0 to 128-size by size
        if  $c_{i+size-1..i+size-1..i+shift} \neq c_{i+size-1..i+shift}^{shift+1}$  then
            raise FixedPointArithmetic
        endif
         $a_{i+size-1..i} \leftarrow c_{i+size-1..i+shift..i} \parallel 0^{shift}$ 
    endfor
X.SHL.I.U.O:
    for i ← 0 to 128-size by size
        if  $c_{i+size-1..i+size..shift} \neq 0^{shift}$  then
            raise FixedPointArithmetic
        endif
         $a_{i+size-1..i} \leftarrow c_{i+size-1..i+shift..i} \parallel 0^{shift}$ 
    endfor
X.ROTR.I:
    for i ← 0 to 128-size by size
         $a_{i+size-1..i} \leftarrow c_{i+shift-1..i} \parallel c_{i+size-1..i+shift}$ 
    endfor
X.SHR.I:
    for i ← 0 to 128-size by size
         $a_{i+size-1..i} \leftarrow c_{i+size-1..i+shift}^{shift} \parallel c_{i+size-1..i+shift}$ 
    endfor
X.SHR.I.U:
    for i ← 0 to 128-size by size
         $a_{i+size-1..i} \leftarrow 0^{shift} \parallel c_{i+size-1..i+shift}$ 
    endfor
endcase
RegWrite(rd, 128, a)
enddef

```

Exceptions

Fixed-point arithmetic
Reserved Instruction

Crossbar Short Immediate Inplace

These operations take operands from two registers and a short immediate value, perform operations on partitions of bits in the operands, and place the concatenated results in the second register.


Operation codes

X.SHL.M.I.2	Crossbar shift left merge immediate pecks
X.SHL.M.I.4	Crossbar shift left merge immediate nibbles
X.SHL.M.I.8	Crossbar shift left merge immediate bytes
X.SHL.M.I.16	Crossbar shift left merge immediate doublets
X.SHL.M.I.32	Crossbar shift left merge immediate quadlets
X.SHL.M.I.64	Crossbar shift left merge immediate octlets
X.SHL.M.I.128	Crossbar shift left merge immediate hexlet
X.SHR.M.I.2	Crossbar shift right merge immediate pecks
X.SHR.M.I.4	Crossbar shift right merge immediate nibbles
X.SHR.M.I.8	Crossbar shift right merge immediate bytes
X.SHR.M.I.16	Crossbar shift right merge immediate doublets
X.SHR.M.I.32	Crossbar shift right merge immediate quadlets
X.SHR.M.I.64	Crossbar shift right merge immediate octlets
X.SHR.M.I.128	Crossbar shift right merge immediate hexlet

Format

X.op.size rd@rc,shift

rd=xopsize(rc,shift)



$t \leftarrow 256 - 2 * size + shift$

$op_{1..0} \leftarrow t_{7..6}$

simm $\leftarrow t_{5..0}$

Description

Two 128-bit values are taken from the contents of registers rd and rc. A third operand is taken from simm. The specified operation is performed, and the result is placed in register rd.

This instruction is undefined and causes a reserved instruction exception if the simm field is greater or equal to the size specified.

Definition

```

def CrossbarShortImmediateInplace(op,rd,rc,simm)
    case {op1..0 || simm} of
        0..127:
            size ← 128
        128..191:
            size ← 64
        192..223:
            size ← 32
        224..239:
            size ← 16
        240..247:
            size ← 8
        248..251:
            size ← 4
        252..253:
            size ← 2,
        254..255:
            raise ReservedInstruction
    endcase
    shift ← (op0 || simm) and (size-1)
    c ← RegRead(rc, 128)
    d ← RegRead(rd, 128)
    for i ← 0 to 128-size by size
        case {op5..2 || 0^2} of
            X.SHR.M.I:
                ai+size-1..i ← ci+shift-1..i || di+size-1..i+shift
            X.SHL.M.I:
                ai+size-1..i ← di+size-1-shift..i || ci+shift-1..i
        endcase
    endfor
    RegWrite(rd, 128, a)
enddef

```

Exceptions

Reserved Instruction

Crossbar Shuffle

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a register.


Operation codes

X.SHUFFLE.4	Crossbar shuffle within pecks
X.SHUFFLE.8	Crossbar shuffle within bytes
X.SHUFFLE.16	Crossbar shuffle within doublets
X.SHUFFLE.32	Crossbar shuffle within quadlets
X.SHUFFLE.64	Crossbar shuffle within octlets
X.SHUFFLE.128	Crossbar shuffle within hexlet
X.SHUFFLE.256	Crossbar shuffle within triclet

Format

X.SHUFFLE.256 rd=rc,rb,v,w,h
X.SHUFFLE.size rd=rcb,v,w

rd=xshuffle256(rc,rb,v,w,h)
rd=xshufflesize(rcb,v,w)



rc \leftarrow rb \leftarrow rcb
x \leftarrow log₂(size)
y \leftarrow log₂(v)
z \leftarrow log₂(w)
op \leftarrow ((x*x*x-3*x*x-4*x)/6-(z*z-z)/2+x*z+y) + (size=256)*(h*32-56)

Description


One of two operations are performed, depending on whether the **rc** and **rb** fields are equal.

If the **rc** and **rb** fields are equal, a 128-bit operand is taken from the contents of register **rc**. Items of size **v** are divided into **w** piles and shuffled together, within groups of size **b**, according to the value of **op**. The result is placed in register **rd**.


If the **rc** and **rb** fields are not equal, the contents of registers **rc** and **rb** are catenated into a 256-bit operand. Items of size **v** are divided into **w** piles and shuffled together, according to the value of **op**. Depending on the value of **h**, a sub-field of **op**, the low 128 bits (**h=0**), or the high 128 bits (**h=1**) of the 256-bit shuffled contents are selected as the result. The result is placed in register **rd**.

This instruction is undefined and causes a reserved instruction exception if **rc** and **rb** are not equal and the **op** field is greater or equal to 56, or if **rc** and **rb** are equal and **op4..0** is greater or equal to 28.

A crossbar 4-way shuffle of bytes within hexlet instruction (X.SHUFFLE.128 rd=rcb,8,4) divides the 128-bit operand into 16 bytes and partitions the bytes 4 ways (indicated by varying shade in the diagram below). The 4 partitions are perfectly shuffled, producing a 128-bit result.



A crossbar 4-way shuffle of bytes within triclet instruction (X.SHUFFLE.256 rd=rc,rb,8,4,0) catenates the contents of **rc** and **rb**, then divides the 256-bit content into 32 bytes and partitions the bytes 4 ways (indicated by varying shade in the diagram below). The low-order halves of the 4 partitions are perfectly shuffled, producing a 128-bit result.



Changing the last immediate value **h** to 1 (X.SHUFFLE.256 rd=rc,rb,8,4,1) modifies the operation to perform the same function on the high-order halves of the 4 partitions.

When rc and rb are equal, the table below shows the value of the **op** field and associated values for size, v, and w.

op	size	v	w
0	4	1	2
1	8	1	2
2	8	2	2
3	8	1	4
4	16	1	2
5	16	2	2
6	16	4	2
7	16	1	4
8	16	2	4
9	16	1	8
10	32	1	2
11	32	2	2
12	32	4	2
13	32	8	2
14	32	1	4
15	32	2	4
16	32	4	4
17	32	1	8
18	32	2	8
19	32	1	16
20	64	1	2
21	64	2	2
22	64	4	2
23	64	8	2
24	64	16	2
25	64	1	4
26	64	2	4
27	64	4	4

op	size	v	w
28	64	8	4
29	64	1	8
30	64	2	8
31	64	4	8
32	64	1	16
33	64	2	16
34	64	1	32
35	128	1	2
36	128	2	2
37	128	4	2
38	128	8	2
39	128	16	2
40	128	32	2
41	128	1	4
42	128	2	4
43	128	4	4
44	128	8	4
45	128	16	4
46	128	1	8
47	128	2	8
48	128	4	8
49	128	8	8
50	128	1	16
51	128	2	16
52	128	4	16
53	128	1	32
54	128	2	32
55	128	1	64

When **rc** and **rb** are not equal, the table below shows the value of the **op4..0** field and associated values for **size**, **v**, and **w**: **Op5** is the value of **h**, which controls whether the low-order or high-order half of each partition is shuffled into the result.

op4..0	size	v	w
0	256	1	2
1	256	2	2
2	256	4	2
3	256	8	2
4	256	16	2
5	256	32	2
6	256	64	2
7	256	1	4
8	256	2	4
9	256	4	4
10	256	8	4
11	256	16	4
12	256	32	4
13	256	1	8
14	256	2	8
15	256	4	8
16	256	8	8
17	256	16	8
18	256	1	16
19	256	2	16
20	256	4	16
21	256	8	16
22	256	1	32
23	256	2	32
24	256	4	32
25	256	1	64
26	256	2	64
27	256	1	128

Definition

```

def CrossbarShuffle(major,rd,rc,rb,op)
  c ← RegRead(rc, 128)
  b ← RegRead(rb, 128)
  if rc=rb then
    case op of
      0..55:
        for x ← 2 to 7; for y ← 0 to x-2; for z ← 1 to x-y-1
          if op =  $\{(x^2x^3x^4x^5x^6)/6 - (z^2z^3z^4z^5z^6)/2 + x^2y^3z^4\}$  then
            for i ← 0 to 127
               $a_i \leftarrow c[i_6..x \parallel i_{y+2..y} \parallel i_{x-1..y+z} \parallel i_{y-1..0}]$ 
            end
          endif
        endfor; endfor; endfor
      56..63:
        raise ReservedInstruction
    end
  end
end

```

```
    endcase
  elseif
    case op4..0 of
      0..27:
        cb ← c || b
        x ← 8
        h ← op5
        for y ← 0 to x-2; for z ← 1 to x-y-1
          if op4..0 = ((17*z-z*z)/2-8+y) then
            for i ← h*128 to 127+h*128
              ai-h*128 ← cb[iy+z-1..y || ix-1..y+z || iy-1..0]
            end
          endif
        endfor; endfor
      28..31:
        raise ReservedInstruction
    endcase
  endif
  RegWrite(rd, 128, a)
enddef
```

Exceptions

Reserved Instruction

Crossbar Swizzle

These operations perform calculations with a general register value and immediate values, placing the result in a general register.


Operation codes

X.SWIZZLE	Crossbar swizzle
-----------	------------------

Format

X.SWIZZLE rd=rc,icopy,iswap

rd=xswizzle(rc,icopy,iswap)



icopya \leftarrow icopy_{5..0}
 iswapa \leftarrow iswap_{5..0}
 ih \leftarrow icopy₆ || iswap₆

Description

The contents of register rc are fetched, and 7-bit immediate values, icopy and iswap, are constructed from the 2-bit ih field and from the 6-bit icopya and iswapa fields. The specified operation is performed on these operands. The result is placed into register rd.

Definition

```
def GroupSwizzleImmediate(ih,rd,rc,icopya,iswapa) as
    icopy  $\leftarrow$  ih1 || icopya
    iswap  $\leftarrow$  ih0 || iswapa
    c  $\leftarrow$  RegRead(rc, 128)
    for i  $\leftarrow$  0 to 127
        ai  $\leftarrow$  c[i] & icopy ^ iswap
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Crossbar Ternary

These operations take three values from registers, perform a group of calculations on partitions of bits of the operands and place the catenated results in a fourth register.


Operation codes

X.SELECT.8	Crossbar select bytes
------------	-----------------------

Format

op ra=rd,rc,rb

ra=op(rd,rc,rb)



Description

The contents of registers rd, rc, and rb are fetched. The specified operation is performed on these operands. The result is placed into register ra.

Definition

```
def CrossbarTernary(op,rd,rc,rb,ra) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    dc ← d || c
    for i ← 0 to 15
        j ← b8*i+4..8*i
        a8*i+7..8*i ← dc8*j+7..8*j
    endfor
    RegWrite(ra, 128, a)
enddef
```

Exceptions

none

Ensemble

These operations take operands from two registers, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

E.CON.8	Ensemble convolve signed bytes
E.CON.16	Ensemble convolve signed doublets
E.CON.32	Ensemble convolve signed quadlets
E.CON.64	Ensemble convolve signed octlets
E.CON.C.8	Ensemble convolve complex bytes
E.CON.C.16	Ensemble convolve complex doublets
E.CON.C.32	Ensemble convolve complex quadlets
E.CON.M.8	Ensemble convolve mixed-signed bytes
E.CON.M.16	Ensemble convolve mixed-signed doublets
E.CON.M.32	Ensemble convolve mixed-signed quadlets
E.CON.M.64	Ensemble convolve mixed-signed octlets
E.CON.U.8	Ensemble convolve unsigned bytes
E.CON.U.16	Ensemble convolve unsigned doublets
E.CON.U.32	Ensemble convolve unsigned quadlets
E.CON.U.64	Ensemble convolve unsigned octlets
E.DIV.64	Ensemble divide signed octlets
E.DIV.U.64	Ensemble divide unsigned octlets
E.MUL.8	Ensemble multiply signed bytes
E.MUL.16	Ensemble multiply signed doublets
E.MUL.32	Ensemble multiply signed quadlets
E.MUL.64	Ensemble multiply signed octlets
E.MUL.SUM.8	Ensemble multiply sum signed bytes
E.MUL.SUM.16	Ensemble multiply sum signed doublets
E.MUL.SUM.32	Ensemble multiply sum signed quadlets
E.MUL.SUM.64	Ensemble multiply sum signed octlets
E.MUL.C.8	Ensemble complex multiply bytes
E.MUL.C.16	Ensemble complex multiply doublets
E.MUL.C.32	Ensemble complex multiply quadlets
E.MUL.M.8	Ensemble multiply mixed-signed bytes
E.MUL.M.16	Ensemble multiply mixed-signed doublets
E.MUL.M.32	Ensemble multiply mixed-signed quadlets
E.MUL.M.64	Ensemble multiply mixed-signed octlets
E.MUL.P.8	Ensemble multiply polynomial bytes
E.MUL.P.16	Ensemble multiply polynomial doublets
E.MUL.P.32	Ensemble multiply polynomial quadlets
E.MUL.P.64	Ensemble multiply polynomial octlets
E.MUL.SUM.C.8	Ensemble multiply sum complex bytes
E.MUL.SUM.C.16	Ensemble multiply sum complex doublets
E.MUL.SUM.C.32	Ensemble multiply sum complex quadlets

E.MUL.SUM.M.8	Ensemble multiply sum mixed-signed bytes
E.MUL.SUM.M.16	Ensemble multiply sum mixed-signed doublets
E.MUL.SUM.M.32	Ensemble multiply sum mixed-signed quadlets
E.MUL.SUM.M.64	Ensemble multiply sum mixed-signed octlets
E.MUL.SUM.U.8	Ensemble multiply sum unsigned bytes
E.MUL.SUM.U.16	Ensemble multiply sum unsigned doublets
E.MUL.SUM.U.32	Ensemble multiply sum unsigned quadlets
E.MUL.SUM.U.64	Ensemble multiply sum unsigned octlets
E.MUL.U.8	Ensemble multiply unsigned bytes
E.MUL.U.16	Ensemble multiply unsigned doublets
E.MUL.U.32	Ensemble multiply unsigned quadlets
E.MUL.U.64	Ensemble multiply unsigned octlets


Selection

class	op	type	size					
multiply	E.MUL	NONE	M	U	P	8	16	32
		C				8	16	32
multiply sum	E.MUL.SUM	NONE	M	U		8	16	32
		C				8	16	32
divide	E.DIV	NONE	U					64

Format

E.op.size rd=rc,rb

rd=eopsize(rc,rb)

Description

Two values are taken from the contents of registers rc and rb. The specified operation is performed, and the result is placed in register rd.

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&Vsize-1+i)h-size || Vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def c ← PolyMultiply(size,a,b) as
    p[0] ← 02*size
    for k ← 0 to size-1
        p[k+1] ← p[k] ^ a_k ? (0size-k || b || 0k) : 02*size
    endfor
    c ← p[size]

```

```
enddef
```

```
def Ensemble(op,size,rd,rc,rb)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case op of
        E.MUL; E.MUL.C; EMULSUM, E.MUL.SUM.C, E.CON, E.CON.C, E.DIV:
            cs ← bs ← 1
        E.MUL.M; EMULSUM.M, E.CON.M:
            cs ← 0
            bs ← 1
        E.MUL.U; EMULSUM.U, E.CON.U, E.DIV.U, E.MUL.P:
            cs ← bs ← 0
    endcase
    case op of
        E.MUL, E.MUL.U, E.MUL.M:
            for i ← 0 to 64-size by size
                d2*(i+size)-1..2*i ← mul(size,2*size,cs,c,i,bs,b,i)
            endfor
        E.MUL.P:
            for i ← 0 to 64-size by size
                d2*(i+size)-1..2*i ← PolyMultiply(size,csize-1+i..i,bsize-1+i..i)
            endfor
        E.MUL.C:
            for i ← 0 to 64-size by size
                if (i and size) = 0 then
                    p ← mul(size,2*size,1,c,i,1,b,i) - mul(size,2*size,1,c,i+size,1,b,i+size)
                else
                    p ← mul(size,2*size,1,c,i,1,b,i+size) + mul(size,2*size,1,c,i,1,b,i+size)
                endif
                d2*(i+size)-1..2*i ← p
            endfor
        E.MUL.SUM, E.MUL.SUM.U, E.MUL.SUM.M:
            p[0] ← 0128
            for i ← 0 to 128-size by size
                p[i+size] ← p[i] + mul(size,128,cs,c,i,bs,b,i)
            endfor
            a ← p[128]
        E.MUL.SUM.C:
            p[0] ← 064
            p[size] ← 064
            for i ← 0 to 128-size by size
                if (i and size) = 0 then
                    p[i+2*size] ← p[i] + mul(size,64,1,c,i,1,b,i)
                    - mul(size,64,1,c,i+size,1,b,i+size)
                else
                    p[i+2*size] ← p[i] + mul(size,64,1,c,i,1,b,i+size)
                    + mul(size,64,1,c,i+size,1,b,i)
                endif
            endfor
            a ← p[128+size] || p[128]
        E.CON, E.CON.U, E.CON.M:
            p[0] ← 0128
            for j ← 0 to 64-size by size
```

```

for i ← 0 to 64-size by size
    p[j+size]2*(i+size)-1..2*i ← p[j]2*(i+size)-1..2*i +
        mul(size,2*size,cs,c,i+64-j,bs,b,j)
endfor
endfor
a ← p[64]
E.CON.C:
p[0] ← 0128
for j ← 0 to 64-size by size
    for i ← 0 to 64-size by size
        if ((-i) and j and size) = 0 then
            p[j+size]2*(i+size)-1..2*i ← p[j]2*(i+size)-1..2*i +
                mul(size,2*size,1,c,i+64-j,1,b,j)
        else
            p[j+size]2*(i+size)-1..2*i ← p[j]2*(i+size)-1..2*i -
                mul(size,2*size,1,c,i+64-j+2*size,1,b,j)
        endif
    endfor
endfor
a ← p[64]
E.DIV:
if (b = 0) or ( (c = (111063)) and (b = 164) ) then
    a ← undefined
else
    q ← c / b
    r ← c - q*b
    a ← r63..0 || q63..0
endif
E.DIV.U:
if b = 0 then
    a ← undefined
else
    q ← (0 || c) / (0 || b)
    r ← c - (0 || q)*(0 || b)
    a ← r63..0 || q63..0
endif
endcase
RegWrite(rd, 128, a)
enddef

```

Exceptions

none

Ensemble Convolve Extract Immediate

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.

Operation codes

E.CON.X.I.C.8.C.B	Ensemble convolve extract immediate signed complex bytes big-endian ceiling
E.CON.X.I.C.8.F.B	Ensemble convolve extract immediate signed complex bytes big-endian floor
E.CON.X.I.C.8.N.B	Ensemble convolve extract immediate signed complex bytes big-endian nearest
E.CON.X.I.C.8.Z.B	Ensemble convolve extract immediate signed complex bytes big-endian zero
E.CON.X.I.C.16.C.B	Ensemble convolve extract immediate signed complex doublets big-endian ceiling
E.CON.X.I.C.16.F.B	Ensemble convolve extract immediate signed complex doublets big-endian floor
E.CON.X.I.C.16.N.B	Ensemble convolve extract immediate signed complex doublets big-endian nearest
E.CON.X.I.C.16.Z.B	Ensemble convolve extract immediate signed complex doublets big-endian zero
E.CON.X.I.C.32.C.B	Ensemble convolve extract immediate signed complex quadlets big-endian ceiling
E.CON.X.I.C.32.F.B	Ensemble convolve extract immediate signed complex quadlets big-endian floor
E.CON.X.I.C.32.N.B	Ensemble convolve extract immediate signed complex quadlets big-endian nearest
E.CON.X.I.C.32.Z.B	Ensemble convolve extract immediate signed complex quadlets big-endian zero
E.CON.X.I.C.64.C.B	Ensemble convolve extract immediate signed complex octlets big-endian ceiling
E.CON.X.I.C.64.F.B	Ensemble convolve extract immediate signed complex octlets big-endian floor
E.CON.X.I.C.64.N.B	Ensemble convolve extract immediate signed complex octlets big-endian nearest
E.CON.X.I.C.64.Z.B	Ensemble convolve extract immediate signed complex octlets big-endian zero
E.CON.X.I.C.8.C.L	Ensemble convolve extract immediate signed complex bytes little-endian ceiling
E.CON.X.I.C.8.F.L	Ensemble convolve extract immediate signed complex bytes little-endian floor
E.CON.X.I.C.8.N.L	Ensemble convolve extract immediate signed complex bytes little-endian nearest
E.CON.X.I.C.8.Z.L	Ensemble convolve extract immediate signed complex bytes little-endian zero
E.CON.X.I.C.16.C.L	Ensemble convolve extract immediate signed complex doublets little-endian ceiling
E.CON.X.I.C.16.F.L	Ensemble convolve extract immediate signed complex doublets little-endian floor
E.CON.X.I.C.16.N.L	Ensemble convolve extract immediate signed complex doublets little-endian nearest
E.CON.X.I.C.16.Z.L	Ensemble convolve extract immediate signed complex doublets little-endian zero
E.CON.X.I.C.32.C.L	Ensemble convolve extract immediate signed complex quadlets little-endian ceiling
E.CON.X.I.C.32.F.L	Ensemble convolve extract immediate signed complex quadlets little-endian floor
E.CON.X.I.C.32.N.L	Ensemble convolve extract immediate signed complex quadlets little-endian nearest
E.CON.X.I.C.32.Z.L	Ensemble convolve extract immediate signed complex quadlets little-endian zero
E.CON.X.I.C.64.C.L	Ensemble convolve extract immediate signed complex octlets little-endian ceiling
E.CON.X.I.C.64.F.L	Ensemble convolve extract immediate signed complex octlets little-endian floor
E.CON.X.I.C.64.N.L	Ensemble convolve extract immediate signed complex octlets little-endian nearest
E.CON.X.I.C.64.Z.L	Ensemble convolve extract immediate signed complex octlets little-endian zero
E.CON.X.I.8.C.B	Ensemble convolve extract immediate signed bytes big-endian ceiling
E.CON.X.I.8.F.B	Ensemble convolve extract immediate signed bytes big-endian floor
E.CON.X.I.8.N.B	Ensemble convolve extract immediate signed bytes big-endian nearest
E.CON.X.I.8.Z.B	Ensemble convolve extract immediate signed bytes big-endian zero
E.CON.X.I.16.C.B	Ensemble convolve extract immediate signed doublets big-endian ceiling

E.CON.X.I.16.F.B	Ensemble convolve extract immediate signed doublets big-endian floor
E.CON.X.I.16.N.B	Ensemble convolve extract immediate signed doublets big-endian nearest
E.CON.X.I.16.Z.B	Ensemble convolve extract immediate signed doublets big-endian zero
E.CON.X.I.32.C.B	Ensemble convolve extract immediate signed quadlets big-endian ceiling
E.CON.X.I.32.F.B	Ensemble convolve extract immediate signed quadlets big-endian floor
E.CON.X.I.32.N.B	Ensemble convolve extract immediate signed quadlets big-endian nearest
E.CON.X.I.32.Z.B	Ensemble convolve extract immediate signed quadlets big-endian zero
E.CON.X.I.64.C.B	Ensemble convolve extract immediate signed octlets big-endian ceiling
E.CON.X.I.64.F.B	Ensemble convolve extract immediate signed octlets big-endian floor
E.CON.X.I.64.N.B	Ensemble convolve extract immediate signed octlets big-endian nearest
E.CON.X.I.64.Z.B	Ensemble convolve extract immediate signed octlets big-endian zero
E.CON.X.I.8.C.L	Ensemble convolve extract immediate signed bytes little-endian ceiling
E.CON.X.I.8.F.L	Ensemble convolve extract immediate signed bytes little-endian floor
E.CON.X.I.8.N.L	Ensemble convolve extract immediate signed bytes little-endian nearest
E.CON.X.I.8.Z.L	Ensemble convolve extract immediate signed bytes little-endian zero
E.CON.X.I.16.C.L	Ensemble convolve extract immediate signed doublets little-endian ceiling
E.CON.X.I.16.F.L	Ensemble convolve extract immediate signed doublets little-endian floor
E.CON.X.I.16.N.L	Ensemble convolve extract immediate signed doublets little-endian nearest
E.CON.X.I.16.Z.L	Ensemble convolve extract immediate signed doublets little-endian zero
E.CON.X.I.32.C.L	Ensemble convolve extract immediate signed quadlets little-endian ceiling
E.CON.X.I.32.F.L	Ensemble convolve extract immediate signed quadlets little-endian floor
E.CON.X.I.32.N.L	Ensemble convolve extract immediate signed quadlets little-endian nearest
E.CON.X.I.32.Z.L	Ensemble convolve extract immediate signed quadlets little-endian zero
E.CON.X.I.64.C.L	Ensemble convolve extract immediate signed octlets little-endian ceiling
E.CON.X.I.64.F.L	Ensemble convolve extract immediate signed octlets little-endian floor
E.CON.X.I.64.N.L	Ensemble convolve extract immediate signed octlets little-endian nearest
E.CON.X.I.64.Z.L	Ensemble convolve extract immediate signed octlets little-endian zero
E.CON.X.I.M.8.C.B	Ensemble convolve extract immediate mixed-signed bytes big-endian ceiling
E.CON.X.I.M.8.F.B	Ensemble convolve extract immediate mixed-signed bytes big-endian floor
E.CON.X.I.M.8.N.B	Ensemble convolve extract immediate mixed-signed bytes big-endian nearest
E.CON.X.I.M.8.Z.B	Ensemble convolve extract immediate mixed-signed bytes big-endian zero
E.CON.X.I.M.16.C.B	Ensemble convolve extract immediate mixed-signed doublets big-endian ceiling
E.CON.X.I.M.16.F.B	Ensemble convolve extract immediate mixed-signed doublets big-endian floor
E.CON.X.I.M.16.N.B	Ensemble convolve extract immediate mixed-signed doublets big-endian nearest
E.CON.X.I.M.16.Z.B	Ensemble convolve extract immediate mixed-signed doublets big-endian zero
E.CON.X.I.M.32.C.B	Ensemble convolve extract immediate mixed-signed quadlets big-endian ceiling
E.CON.X.I.M.32.F.B	Ensemble convolve extract immediate mixed-signed quadlets big-endian floor
E.CON.X.I.M.32.N.B	Ensemble convolve extract immediate mixed-signed quadlets big-endian nearest
E.CON.X.I.M.32.Z.B	Ensemble convolve extract immediate mixed-signed quadlets big-endian zero
E.CON.X.I.M.64.C.B	Ensemble convolve extract immediate mixed-signed octlets big-endian ceiling
E.CON.X.I.M.64.F.B	Ensemble convolve extract immediate mixed-signed octlets big-endian floor
E.CON.X.I.M.64.N.B	Ensemble convolve extract immediate mixed-signed octlets big-endian nearest
E.CON.X.I.M.64.Z.B	Ensemble convolve extract immediate mixed-signed octlets big-endian zero
E.CON.X.I.M.8.C.L	Ensemble convolve extract immediate mixed-signed bytes little-endian ceiling
E.CON.X.I.M.8.F.L	Ensemble convolve extract immediate mixed-signed bytes little-endian floor
E.CON.X.I.M.8.N.L	Ensemble convolve extract immediate mixed-signed bytes little-endian nearest

E.CONX.I.M.8.Z.L	Ensemble convolve extract immediate mixed-signed bytes little-endian zero
E.CONX.I.M.16.C.L	Ensemble convolve extract immediate mixed-signed doublets little-endian ceiling
E.CONX.I.M.16.F.L	Ensemble convolve extract immediate mixed-signed doublets little-endian floor
E.CONX.I.M.16.N.L	Ensemble convolve extract immediate mixed-signed doublets little-endian nearest
E.CONX.I.M.16.Z.L	Ensemble convolve extract immediate mixed-signed doublets little-endian zero
E.CONX.I.M.32.C.L	Ensemble convolve extract immediate mixed-signed quadlets little-endian ceiling
E.CONX.I.M.32.F.L	Ensemble convolve extract immediate mixed-signed quadlets little-endian floor
E.CONX.I.M.32.N.L	Ensemble convolve extract immediate mixed-signed quadlets little-endian nearest
E.CONX.I.M.32.Z.L	Ensemble convolve extract immediate mixed-signed quadlets little-endian zero
E.CONX.I.M.64.C.L	Ensemble convolve extract immediate mixed-signed octlets little-endian ceiling
E.CONX.I.M.64.F.L	Ensemble convolve extract immediate mixed-signed octlets little-endian floor
E.CONX.I.M.64.N.L	Ensemble convolve extract immediate mixed-signed octlets little-endian nearest
E.CONX.I.M.64.Z.L	Ensemble convolve extract immediate mixed-signed octlets little-endian zero
E.CONX.I.U.8.C.B	Ensemble convolve extract immediate unsigned bytes big-endian ceiling
E.CONX.I.U.8.F.B	Ensemble convolve extract immediate unsigned bytes big-endian floor
E.CONX.I.U.8.N.B	Ensemble convolve extract immediate unsigned bytes big-endian nearest
E.CONX.I.U.16.C.B	Ensemble convolve extract immediate unsigned doublets big-endian ceiling
E.CONX.I.U.16.F.B	Ensemble convolve extract immediate unsigned doublets big-endian floor
E.CONX.I.U.16.N.B	Ensemble convolve extract immediate unsigned doublets big-endian nearest
E.CONX.I.U.32.C.B	Ensemble convolve extract immediate unsigned quadlets big-endian ceiling
E.CONX.I.U.32.F.B	Ensemble convolve extract immediate unsigned quadlets big-endian floor
E.CONX.I.U.32.N.B	Ensemble convolve extract immediate unsigned quadlets big-endian nearest
E.CONX.I.U.64.C.B	Ensemble convolve extract immediate unsigned octlets big-endian ceiling
E.CONX.I.U.64.F.B	Ensemble convolve extract immediate unsigned octlets big-endian floor
E.CONX.I.U.64.N.B	Ensemble convolve extract immediate unsigned octlets big-endian nearest
E.CONX.I.U.8.C.L	Ensemble convolve extract immediate unsigned bytes little-endian ceiling
E.CONX.I.U.8.F.L	Ensemble convolve extract immediate unsigned bytes little-endian floor
E.CONX.I.U.8.N.L	Ensemble convolve extract immediate unsigned bytes little-endian nearest
E.CONX.I.U.16.C.L	Ensemble convolve extract immediate unsigned doublets little-endian ceiling
E.CONX.I.U.16.F.L	Ensemble convolve extract immediate unsigned doublets little-endian floor
E.CONX.I.U.16.N.L	Ensemble convolve extract immediate unsigned doublets little-endian nearest
E.CONX.I.U.32.C.L	Ensemble convolve extract immediate unsigned quadlets little-endian ceiling
E.CONX.I.U.32.F.L	Ensemble convolve extract immediate unsigned quadlets little-endian floor
E.CONX.I.U.32.N.L	Ensemble convolve extract immediate unsigned quadlets little-endian nearest
E.CONX.I.U.64.C.L	Ensemble convolve extract immediate unsigned octlets little-endian ceiling
E.CONX.I.U.64.F.L	Ensemble convolve extract immediate unsigned octlets little-endian floor
E.CONX.I.U.64.N.L	Ensemble convolve extract immediate unsigned octlets little-endian nearest

Format

E.op.size.rnd rd@rc,rb,i

rd=eopsizernd(rd,rc,rb,i)

31	24 23	18 17	12 11	6 5 4 3 2 1 0
E.op		rd	rc	rb sz rnd sh
8	6	6	6	2 2 2

sz $\leftarrow \log(\text{size}) - 3$

sh $\leftarrow \text{size} + 7 - \log(\text{size}) - i$


Description

The contents of registers rd and rc are catenated, as specified by the order parameter, and used as a first value. A second value is the contents of register rb. The values are partitioned into groups of operands of the size specified and are convolved, producing a group of values. The group of values is rounded, and limited as specified, yielding a group of results which is the size specified. The group of results is catenated and placed in register rd.

Z (zero) rounding is not defined for unsigned extract operations, and a ReservedInstruction exception is raised if attempted. F (floor) rounding will properly round unsigned results downward.


The order parameter of the instruction specifies the order in which the contents of registers rd and rc are catenated. The choice is significant because the contents of register rd is overwritten. When little-endian order is specified, the contents are catenated so that the contents of register rc is most significant (left) and the contents of register rd is least significant (right). When big-endian order is specified, the contents are catenated so that the contents of register rd is most significant (left) and the contents of register rc is least significant (right).

An ensemble-convolve-extract-immediate-doublets instruction (ECON.X.I16, ECON.X.IM16, or ECON.X.IU16) convolves vector [x w v u t s r q p o n m l k j i] with vector [h g f e d c b a], yielding the products [ax+bx+cx+dx+ex+fx+gx+ ... ar+bq+cp+do+en+fm+gl+hk aq+bp+co+dn+em+fl+gk+hj], rounded and limited as specified:



Ensemble convolve extract immediate doublets

An ensemble-convolve-extract-immediate-complex-doubles instruction (ECON.X.IC16) convolves vector [x w v u t s r q p o n m l k j i] with vector [h g f e d c b a], yielding the products [ax+bx+cv+du+et+fs+gr+hq ... as-bt+cq-dr+eo-fp+gm-hn ar+bq+cp+do+en+fm+gl+hk aq-br+co-dp+em-fn+gk+hl], rounded and limited as specified:



Ensemble convolve extract immediate complex doublets

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def EnsembleConvolveExtractImmediate(op,rnd,gsize,rd,rc,rb,sh)
    d ← RegRead(rd, 128)
    c ← RegRead(rd, 128)
    b ← RegRead(rb, 128)
    lgsize ← log(gsize)
    wsize ← 128
    msize ← 256
    vsize ← 128
    case op of
        E.CON.X.I.B, E.CON.X.I.U.B, E.CON.X.I.M.B, E.CON.X.I.C.B:
            m ← d || c
        E.CON.X.I.L, E.CON.X.I.U.L, E.CON.X.I.M.L, E.CON.X.I.C.L:
            m ← c || d
    endcase
    case op of
        E.CON.X.I.U.B, E.CON.X.I.U.L:
            as ← ms ← bs ← false
        E.CON.X.I.M.B, E.CON.X.I.M.L:
    endcase

```

```

        ms ← false
        as ← bs ← true
E.CON.X.I.B, E.CON.X.I.L, E.CON.X.I.C.B, E.CON.X.I.C.L:
        as ← ms ← bs ← true
endcase
h ← {2*gsize} + 7 - lgsiz
r ← h - size - sh
for i ← 0 to wsize-gsize by gsize
    q[0] ← 02*gsize+7-lgsiz
    for j ← 0 to vsize-gsize by gsize
        case op of
            E.CON.X.I.B, E.CON.X.I.L, E.CON.X.I.M.B, E.CON.X.I.M.L,
            E.CON.X.I.U.B, E.CON.X.I.U.L:
                q[j+gsize] ← q[j] + mul(gsize,h,ms,m,i+128-j,bs,b,j)
            E.CON.X.I.C.B, E.CON.X.I.C.L:
                if (-i) & j & gsize = 0 then
                    q[j+gsize] ← q[j] + mul(gsize,h,ms,m,i+128-j,bs,b,j)
                else
                    q[j+gsize] ← q[j] - mul(gsize,h,ms,m,i+128-j+2*gsize,bs,b,j)
                endif
            endcase
        endfor
        p ← q[vsize]
        case rnd of
            none, N:
                s ← 0h-r || -pr || pr-1
            Z:
                s ← 0h-r || pr-1
            F:
                s ← 0h
            C:
                s ← 0h-r || 1^r
        endcase
        v ← ((as & ph-1) || p) + (011s)
        if (vh..r+gsize = (as & vr+gsize-1)h+1-r-gsize then
            agsize-1+i..i ← vgsize-1+r..r
        else
            agsize-1+i..i ← as ? (vh || -vgsize-1) : 1gsize
        endif
    endfor
    a127..wsize ← 0
    RegWrite(rd, 128, a)
enddef

```

Exceptions

none

Ensemble Convolve Floating-point

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register .


Operation codes

E.CON.F.16.B	Ensemble convolve floating-point half big-endian
E.CON.F.16.L	Ensemble convolve floating-point half little-endian
E.CON.F.32.B	Ensemble convolve floating-point single big-endian
E.CON.F.32.L	Ensemble convolve floating-point single little-endian
E.CON.F.64.B	Ensemble convolve floating-point double big-endian
E.CON.F.64.L	Ensemble convolve floating-point double little-endian
E.CON.C.F.16.B	Ensemble convolve complex floating-point half big-endian
E.CON.C.F.16.L	Ensemble convolve complex floating-point half little-endian
E.CON.C.F.32.B	Ensemble convolve complex floating-point single big-endian
E.CON.C.F.32.L	Ensemble convolve complex floating-point single little-endian
E.CON.C.F.64.B	Ensemble convolve complex floating-point double big-endian
E.CON.C.F.64.L	Ensemble convolve complex floating-point double little-endian

Format

E.op.size.order rd=rc,rb


rd=eopsizeorder(rd,rc,rb)



Description


The first value is the catenation of the contents of register rd and rc, as specified by the order parameter. A second value is the contents of register rb. The values are partitioned into groups of operands of the size specified. The second values are multiplied with the first values, then summed, producing a group of result values. The group of result values is catenated and placed in register rd.

An ensemble-convolve-floating-point-half-little-endian instruction (E.CON.F.16.L) convolves vector [x w v u t s r q p o n m l k j i] with vector [h g f e d c b a], yielding the products [ax+bx+cx+...+hx] ... [as+bs+cs+...+hs] [aq+bp+co+dn+em+fl+gk+hj]:



Ensemble convolve floating-point half little-endian

A ensemble-convolve-complex-floating-point-half-little-endian instruction (E.CON.C.F.16.L) convolves vector [x w v u t s r q p o n m l k j i] with vector [h g f e d c b a], yielding the products [ax+bw+cv+du+et+fs+gr+hq ... as-bt+cq-dr+eo-fp+gm-hn ar+bq+cp+do+en+fm+gl+hk aq-br+co-dp+em-fn+gk+hl]:



Ensemble convolve complex floating-point half little-endian

Definition

```

def mul|size,v,i,w,j| as
    mul ← fmul(F(size,vsize-1+i..i),F(size,wsize-1+j..j))
enddef

def EnsembleConvolveFloatingPoint(op,gsize,rd,rc,rb)
    d ← RegRead[rd, 128]
    c ← RegRead[rc, 128]
    b ← RegRead[rb, 128]
    lgsiz ← log(gsize)
    wsize ← 128
    msiz ← 256
    vsize ← 128
    case op of
        E.CON.F.B, E.CON.C.F.B:
            m ← d || c
        E.CON.F.L, E.CON.C.F.L :
            m ← c || d
    endcase
    for i ← 0 to wsize-gsize by gsize
        //NULL value doesn't combine with zero to alter sign bit
        q[0].t ← NULL
        for j ← 0 to vsize-gsize by gsize
            case op of
                E.CON.F.L, E.CON.F.B:

```

```
a[j+gsize] ← fadd[q[j], mul(gsize,m,i+128-j,b,j)]
E.CONCF.L, E.CONCF.B:
    if (~i) & j & gsize = 0 then
        q[j+gsize] ← fadd[q[j], mul(gsize,m,i+128-j,b,j)]
    else
        q[j+gsize] ← fsub[q[j], mul(gsize,m,i+128-j+2*gsize,b,j)]
    endif
endcase
endfor
a[gsize-1+i..i] ← PackF[gsize,q[vsize],N]
endfor
a[127..wsize] ← 0
RegWrite(rd, 128, a)
enddef
```

Exceptions

none

Ensemble Extract

These operations take operands from three registers, perform operations on partitions of bits in the operands, and place the concatenated results in a fourth register.


Operation codes

E.MULX	Ensemble multiply extract
E.EXTRACT	Ensemble extract
E.SCALADD.X	Ensemble scale add extract

Format

E.opra=rd,rc,rb


ra=gop(rd,rc,rb)



Description

The contents of registers rd, rc, and rb are fetched. The specified operation is performed on these operands. The result is placed into register ra.

Bits 31..0 of the contents of register rb specifies several parameters which control the manner in which data is extracted, and for certain operations, the manner in which the operation is performed. The position of the control fields allows for the source position to be added to a fixed control value for dynamic computation, and allows for the lower 16 bits of the control field to be set for some of the simpler extract cases by a single GCOPYI.128 instruction. The control fields are further arranged so that if only the low order 8 bits are non-zero, a 128-bit extraction with truncation and no rounding is performed.



The table below describes the meaning of each label:

label	bits	meaning
fsize	8	field size
dpos	8	destination position
x	1	reserved
s	1	signed vs. unsigned
n	1	complex vs. real multiplication
m	1	merge vs. extract or mixed-sign vs. same-sign multiplication
l	1	limit: saturation vs. truncation
rnd	2	rounding
gssp	9	group size and source position


The 9-bit gssp field encodes both the group size, gsize, and source position, spos, according to the formula $\text{gssp} = 512 \cdot 4^{\text{gsize}} + \text{spos}$. The group size, gsize, is a power of two in the range 1..128. The source position, spos, is in the range 0..(2^{gsize})-1.

The values in the s, n, m, l, and rnd fields have the following meaning:


values	s	n	m	l	rnd
0	unsigned		extract/same-sign	truncate	F
1	signed	real complex	merge/mixed-sign	saturate	Z
2					N
3					C

For the ESCAL.ADD.X instruction, bits 127..64 of the contents of register rc specifies the multipliers for the multiplicands in registers ra and rb. Specifically, bits $64+2^{\text{gsize}}-1:64+\text{gsize}$ is the multiplier for the contents of register ra, and bits $64+\text{gsize}-1:64$ is the multiplier for the contents of register rb.

An ensemble-multiply-extract-doublets instruction (E.MULX) multiplies vector ra [h g f e d c b a] with vector rb [p o n m l k j i], yielding the result [hp go fn em dl ck bj ai], rounded and limited as specified by rc31..0.




An ensemble-multiply-extract-doublets-complex instruction (E.MUL.X with n set) multiplies operand [h g f e d c b a] by operand [p o n m l k j i], yielding the result [gp+ho go-hp en+fm em-fn cl+dk ck-dl aj+bi ai-bj], rounded and limited as specified. Note that this instruction prefers an organization of complex numbers in which the real part is located to the right (lower precision) of the imaginary part.:



Ensemble complex multiply extract doublets


An ensemble-scale-add-extract-doublets instruction (E.SCAL.ADD.X) multiplies vector ra [h g f e d c b a] with rc95..80 [r] and adds the product to the product of vector rb [p o n m l k j i] with rc79..64 [q], yielding the result [hr+pq gr+oq fr+nq er+mq dr+lq cr+kq br+jq ar+iq], rounded and limited as specified by rc31..0.



Ensemble scale add extract doublets


An ensemble-scale-add-extract-doublets-complex instruction (E.SCLADD.X with n set) multiplies vector ra [h g f e d c b a] with rc127..96 [t s] and adds the product to the product of vector rb [p o n m l k j i] with rc95..64 [r q], yielding the result [hs+gt+pq+or gs-ht+oq-pr]

$fs+et+nq+mr\ es-ft+mq-nr\ ds+ct+lq+kr\ cs-dt+kq-lr\ bs+at+jq+ir\ as-bt+iq-jr]$, rounded and limited as specified by rc31..0.




Ensemble complex scale add extract doublets

For the E.EXTRACT instruction, when $m=0$, the parameters are interpreted to select a fields from the catenated contents of registers **rd** and **rc**, extracting values which are catenated and placed in register **ra**:



Ensemble extract

For an ensemble-merge-extract (G.X when m=1), the parameters are interpreted to merge a fields from the contents of register rd with the contents of register rc. The results are catenated and placed in register ra.



Ensemble merge extract

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
endef

def EnsembleExtract(op,ra,rb,rc,rd) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case b8..0 of
        0..255:
            sgszie ← 128
        256..383:
            sgszie ← 64
        384..447:
            sgszie ← 32
        448..479:
            sgszie ← 16
        480..495:
            sgszie ← 8
        496..503:
            sgszie ← 4
        504..507:
            sgszie ← 2
        508..511:
            sgszie ← 1
    endcase
    l ← b11
    m ← b12
    n ← b13
    signed ← b14

```

```

case op of
  E.EXTRACT:
    gsize ← sgszie
    h ← (2-m)*gsize
    as ← signed
    spos ← (b8..0) and ((2-m)*gsize-1)
  E.SCALADD.X:
    if (sgsize < 8) then
      gsize ← 8
    elseif (sgsize*(n+1) > 32) then
      gsize ← 32/(n+1)
    else
      gsize ← sgszie
    endif
    ds ← cs ← signed
    bs ← signed ^ m
    as ← signed or m or n
    h ← (2*gsize) + 1 + n
    spos ← (b8..0) and (2*gsize-1)
  E.MULX:
    if (sgsize < 8) then
      gsize ← 8
    elseif (sgsize*(n+1) > 128) then
      gsize ← 128/(n+1)
    else
      gsize ← sgszie
    endif
    ds ← signed
    cs ← signed ^ m
    as ← signed or m or n
    h ← (2*gsize) + n
    spos ← (b8..0) and (2*gsize-1)
  endcase
  dpos ← (0 || b23..16) and (gsize-1)
  r ← spos
  sfszie ← (0 || b31..24) and (gsize-1)
  tfszie ← (sfszie = 0) or ((sfszie+dpos) > gsize) ? gsize-dpos : sfszie
  fsize ← (tfszie + spos > h) ? h - spos : tfszie
  if (b10..9 = Z) and not as then
    rnd ← F
  else
    rnd ← b10..9
  endif
  for i ← 0 to 128-gsize by gsize
    case op of
      E.EXTRACT:
        if m then
          p ← dgszie+i..i
        else
          p ← (d || c)2*(gsize+i)-1..2*i
        endif
      E.MULX:
        if n then

```

```

        if (i and gsize) = 0 then
            p ← mul(gsize,h,ds,d,i,cs,c,i) - mul(gsize,h,ds,d,i+size,cs,c,i+size)
        else
            p ← mul(gsize,h,ds,d,i,cs,c,i+size) + mul(gsize,h,ds,d,i,cs,c,i+size)
        endif
    else
        p ← mul(gsize,h,ds,d,i,cs,c,i)
    endif
E.SCALADDX:
    if n then
        if (i and gsize) = 0 then
            p ← mul(gsize,h,ds,d,i,bs,b,64+2*gsize)
            + mul(gsize,h,cs,c,i,bs,b,64)
            - mul(gsize,h,ds,d,i+gsize,bs,b,64+3*gsize)
            - mul(gsize,h,cs,c,i+gsize,bs,b,64+gsize)
        else
            p ← mul(gsize,h,ds,d,i,bs,b,64+3*gsize)
            + mul(gsize,h,cs,c,i,bs,b,64+gsize)
            + mul(gsize,h,ds,d,i+gsize,bs,b,64+2*gsize)
            + mul(gsize,h,cs,c,i+gsize,bs,b,64)
        endif
    else
        p ← mul(gsize,h,ds,d,i,bs,b,64+gsize) + mul(gsize,h,cs,c,i,bs,b,64)
    endif
endcase
case rnd of
    N:
        s ← 0h-r || ~pr || prr-1
    Z:
        s ← 0h-r || prr-1
    F:
        s ← 0h
    C:
        s ← 0h-r || 1r
endcase
v ← ((as & ph-1) || p) + (0 || s)
if (Vh..r+fsize = (as & vr+fsize-1)h+1-r-fsize) or not (l and (op = E.EXTRACT)) then
    w ← (as & vr+fsize-1)gsize-fsize-dpos || vsize-1+r..r || 0dpos
else
    w ← (s ? (vh || ~vgsize-dpos-1) : 1gsize-dpos) || 0dpos
endif
if m and (op = E.EXTRACT) then
    asize-1+i..i ← cgsize-1+i..dpos+fsize+i || wdpos+fsize-1..dpos || cdpos-1+1..i
else
    asize-1+i..i ← w
endif
endfor
RegWrite(ra, 128, a)
enddef

```

Exceptions

none

Ensemble Extract Immediate

These operations take operands from two registers and a short immediate value, perform operations on partitions of bits in the operands, and place the concatenated results in a third register.

Operation codes

E.EXTRACT.I.8.C	Ensemble extract immediate signed bytes ceiling
E.EXTRACT.I.8.F	Ensemble extract immediate signed bytes floor
E.EXTRACT.I.8.N	Ensemble extract immediate signed bytes nearest
E.EXTRACT.I.8.Z	Ensemble extract immediate signed bytes zero
E.EXTRACT.I.16.C	Ensemble extract immediate signed doublets ceiling
E.EXTRACT.I.16.F	Ensemble extract immediate signed doublets floor
E.EXTRACT.I.16.N	Ensemble extract immediate signed doublets nearest
E.EXTRACT.I.16.Z	Ensemble extract immediate signed doublets zero
E.EXTRACT.I.32.C	Ensemble extract immediate signed quadlets ceiling
E.EXTRACT.I.32.F	Ensemble extract immediate signed quadlets floor
E.EXTRACT.I.32.N	Ensemble extract immediate signed quadlets nearest
E.EXTRACT.I.32.Z	Ensemble extract immediate signed quadlets zero
E.EXTRACT.I.64.C	Ensemble extract immediate signed octlets ceiling
E.EXTRACT.I.64.F	Ensemble extract immediate signed octlets floor
E.EXTRACT.I.64.N	Ensemble extract immediate signed octlets nearest
E.EXTRACT.I.64.Z	Ensemble extract immediate signed octlets zero
E.EXTRACT.I.U.8.C	Ensemble extract immediate unsigned bytes ceiling
E.EXTRACT.I.U.8.F	Ensemble extract immediate unsigned bytes floor
E.EXTRACT.I.U.8.N	Ensemble extract immediate unsigned bytes nearest
E.EXTRACT.I.U.16.C	Ensemble extract immediate unsigned doublets ceiling
E.EXTRACT.I.U.16.F	Ensemble extract immediate unsigned doublets floor
E.EXTRACT.I.U.16.N	Ensemble extract immediate unsigned doublets nearest
E.EXTRACT.I.U.32.C	Ensemble extract immediate unsigned quadlets ceiling
E.EXTRACT.I.U.32.F	Ensemble extract immediate unsigned quadlets floor
E.EXTRACT.I.U.32.N	Ensemble extract immediate unsigned quadlets nearest
E.EXTRACT.I.U.64.C	Ensemble extract immediate unsigned octlets ceiling
E.EXTRACT.I.U.64.F	Ensemble extract immediate unsigned octlets floor
E.EXTRACT.I.U.64.N	Ensemble extract immediate unsigned octlets nearest
E.MULX.I.8.C	Ensemble multiply extract immediate signed bytes ceiling
E.MULX.I.8.F	Ensemble multiply extract immediate signed bytes floor
E.MULX.I.8.N	Ensemble multiply extract immediate signed bytes nearest
E.MULX.I.8.Z	Ensemble multiply extract immediate signed bytes zero
E.MULX.I.16.C	Ensemble multiply extract immediate signed doublets ceiling
E.MULX.I.16.F	Ensemble multiply extract immediate signed doublets floor
E.MULX.I.16.N	Ensemble multiply extract immediate signed doublets nearest
E.MULX.I.16.Z	Ensemble multiply extract immediate signed doublets zero
E.MULX.I.32.C	Ensemble multiply extract immediate signed quadlets ceiling
E.MULX.I.32.F	Ensemble multiply extract immediate signed quadlets floor


E.MUL.X.I.32.N	Ensemble multiply extract immediate signed quadlets nearest
E.MUL.X.I.32.Z	Ensemble multiply extract immediate signed quadlets zero
E.MUL.X.I.64.C	Ensemble multiply extract immediate signed octlets ceiling
E.MUL.X.I.64.F	Ensemble multiply extract immediate signed octlets floor
E.MUL.X.I.64.N	Ensemble multiply extract immediate signed octlets nearest
E.MUL.X.I.64.Z	Ensemble multiply extract immediate signed octlets zero
E.MUL.X.I.C.8.C	Ensemble multiply extract immediate complex bytes ceiling
E.MUL.X.I.C.8.F	Ensemble multiply extract immediate complex bytes floor
E.MUL.X.I.C.8.N	Ensemble multiply extract immediate complex bytes nearest
E.MUL.X.I.C.8.Z	Ensemble multiply extract immediate complex bytes zero
E.MUL.X.I.C.16.C	Ensemble multiply extract immediate complex doublets ceiling
E.MUL.X.I.C.16.F	Ensemble multiply extract immediate complex doublets floor
E.MUL.X.I.C.16.N	Ensemble multiply extract immediate complex doublets nearest
E.MUL.X.I.C.16.Z	Ensemble multiply extract immediate complex doublets zero
E.MUL.X.I.C.32.C	Ensemble multiply extract immediate complex quadlets ceiling
E.MUL.X.I.C.32.F	Ensemble multiply extract immediate complex quadlets floor
E.MUL.X.I.C.32.N	Ensemble multiply extract immediate complex quadlets nearest
E.MUL.X.I.C.32.Z	Ensemble multiply extract immediate complex quadlets zero
E.MUL.X.I.C.64.C	Ensemble multiply extract immediate complex octlets ceiling
E.MUL.X.I.C.64.F	Ensemble multiply extract immediate complex octlets floor
E.MUL.X.I.C.64.N	Ensemble multiply extract immediate complex octlets nearest
E.MUL.X.I.C.64.Z	Ensemble multiply extract immediate complex octlets zero
E.MUL.X.I.M.8.C	Ensemble multiply extract immediate mixed-signed bytes ceiling
E.MUL.X.I.M.8.F	Ensemble multiply extract immediate mixed-signed bytes floor
E.MUL.X.I.M.8.N	Ensemble multiply extract immediate mixed-signed bytes nearest
E.MUL.X.I.M.8.Z	Ensemble multiply extract immediate mixed-signed bytes zero
E.MUL.X.I.M.16.C	Ensemble multiply extract immediate mixed-signed doublets ceiling
E.MUL.X.I.M.16.F	Ensemble multiply extract immediate mixed-signed doublets floor
E.MUL.X.I.M.16.N	Ensemble multiply extract immediate mixed-signed doublets nearest
E.MUL.X.I.M.16.Z	Ensemble multiply extract immediate mixed-signed doublets zero
E.MUL.X.I.M.32.C	Ensemble multiply extract immediate mixed-signed quadlets ceiling
E.MUL.X.I.M.32.F	Ensemble multiply extract immediate mixed-signed quadlets floor
E.MUL.X.I.M.32.N	Ensemble multiply extract immediate mixed-signed quadlets nearest
E.MUL.X.I.M.32.Z	Ensemble multiply extract immediate mixed-signed quadlets zero
E.MUL.X.I.M.64.C	Ensemble multiply extract immediate mixed-signed octlets ceiling
E.MUL.X.I.M.64.F	Ensemble multiply extract immediate mixed-signed octlets floor
E.MUL.X.I.M.64.N	Ensemble multiply extract immediate mixed-signed octlets nearest
E.MUL.X.I.M.64.Z	Ensemble multiply extract immediate mixed-signed octlets zero
E.MUL.X.I.U.8.C	Ensemble multiply extract immediate unsigned bytes ceiling
E.MUL.X.I.U.8.F	Ensemble multiply extract immediate unsigned bytes floor
E.MUL.X.I.U.8.N	Ensemble multiply extract immediate unsigned bytes nearest
E.MUL.X.I.U.16.C	Ensemble multiply extract immediate unsigned doublets ceiling
E.MUL.X.I.U.16.F	Ensemble multiply extract immediate unsigned doublets floor
E.MUL.X.I.U.16.N	Ensemble multiply extract immediate unsigned doublets nearest
E.MUL.X.I.U.32.C	Ensemble multiply extract immediate unsigned quadlets ceiling
E.MUL.X.I.U.32.F	Ensemble multiply extract immediate unsigned quadlets floor

E.MUL.X.I.U.32.N	Ensemble multiply extract immediate unsigned quadlets nearest
E.MUL.X.I.U.64.C	Ensemble multiply extract immediate unsigned octlets ceiling
E.MUL.X.I.U.64.F	Ensemble multiply extract immediate unsigned octlets floor
E.MUL.X.I.U.64.N	Ensemble multiply extract immediate unsigned octlets nearest

Format

E.op.size.rnd rd=rc,rb,i

rd=eopsizernd(rc,rb,i)

sz $\leftarrow \log(\text{size}) - 3$

case op of

E.EXTRACT.I, E.EXTRACT.I.U, E.MUL.X.I, E.MULX.I.U, E.MULX.I.M:

 assert size $\geq i \geq \text{size}-3$ sh $\leftarrow \text{size} - i$

E.MUL.X.I.C:

 assert size+1 $\geq i \geq \text{size}-2$ sh $\leftarrow \text{size} + 1 - i$

endcase


Description

The contents of registers rc and rb are partitioned into groups of operands of the size specified and multiplied, added or subtracted, or are catenated and partitioned into operands of twice the size specified. The group of values is rounded, and limited as specified, yielding a group of results, each of which is the size specified. The group of results is catenated and placed in register rd.


For mixed-signed multiplies, the contents of register rc is signed, and the contents of register rb is unsigned. The extraction operation and the result of mixed-signed multiplies is signed.

Z (zero) rounding is not defined for unsigned extract operations, and a ReservedInstruction exception is raised if attempted. F (floor) rounding will properly round unsigned results downward.


An ensemble multiply extract immediate doublets instruction (E.MULXI.16 or E.MUL.X.I.U.16) multiplies operand [h g f e d c b a] by operand [p o n m l k j i], yielding the products [hp go fn em dl ck bj ai], rounded and limited as specified:



Another illustration of ensemble multiply extract immediate doublets instruction (E.MUL.X.I.16 or E.MUL.X.I.U.16):




An ensemble multiply extract immediate complex doublets instruction (E.MULXIC.16 or E.MUL.X.I.U.16) multiplies operand [h g f e d c b a] by operand [p o n m l k j i], yielding the result [gp+ho go-hp en+fm em-fn cl+dk ck-dl aj+bi ai-bj], rounded and limited as specified. Note that this instruction prefers an organization of complex numbers in which the real part is located to the right (lower precision) of the imaginary part.:.



Ensemble multiply extract immediate complex doublets

Another illustration of ensemble multiply extract immediate complex doublets instruction (E.MUL.X.I.C.16 or E.MUL.X.I.U.16).:



Ensemble multiply extract immediate complex doublets

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def EnsembleExtractImmediate(op,rnd,size,ra,rb,rc,sh)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case op of
        E.EXTRACT.I, E.MULX.I, E.MULX.I.C:
            as ← 1
            cs ← 1
            bs ← 1
        E.MULX.I.M:
            as ← 1
            cs ← 0
            bs ← 1
        E.EXTRACT.I.U, E.MULX.I.U:
            as ← 1
            cs ← 0
            bs ← 0
            if rnd = Z then
                raise ReservedInstruction
            endif
    endcase
    case op of
        E.EXTRACT.I, E.EXTRACT.I.U, E.MULX.I, E.MULX.I.U, E.MULX.I.M:
            h ← 2*size
        E.MULX.I.C:
            h ← (2*size) + 1
    endcase
    r ← h - size - sh
    for i ← 0 to 128-size by size
        case op of
            E.EXTRACT.I, E.EXTRACT.I.U:
                p ← (c || b)z(size+i)-1..2^i
            E.MULX.I, E.MULX.I.M, E.MULX.I.U:
                p ← mul(size,h,cs,c,i,bs,b,i)
            E.MULX.I.C:
                if i & size = 0 then
                    p ← mul(size,h,cs,c,i,bs,b,i) - mul(size,h,cs,c,i+size,bs,b,i+size)
                else
                    p ← mul(size,h,cs,c,i,bs,b,i+size) + mul(size,h,cs,c,i,bs,b,i+size)
                endif
        endcase
        case rnd of
            none, N:
                s ← 0h-r || ~pr || prr-1
            Z:
                s ← 0h-r || prr-1
            F:
                s ← 0h
            C:

```

```
s ← 0h-r || 1r
endcase
v ← ((as & ph-1) || p) + (0 || s)
if (vh..r+size = (as & vr+size-1)h+1-r-size then
    asize-1+i..i ← vsize-1+r..r
else
    asize-1+i..i ← as ? (vh || -vhsize-1) : 1size
endif
endfor
RegWrite(rd, 128, a)
enddef
```

Exceptions

ReservedInstruction

Ensemble Extract Immediate Inplace

These operations take operands from two registers and a short immediate value, perform operations on partitions of bits in the operands, and place the catenated results in a third register.

Operation codes

E.MULADD.X.I.C.8.C	Ensemble multiply add extract immediate signed complex bytes ceiling
E.MULADD.X.I.C.8.F	Ensemble multiply add extract immediate signed complex bytes floor
E.MULADD.X.I.C.8.N	Ensemble multiply add extract immediate signed complex bytes nearest
E.MULADD.X.I.C.8.Z	Ensemble multiply add extract immediate signed complex bytes zero
E.MULADD.X.I.C.16.C	Ensemble multiply add extract immediate signed complex doublets ceiling
E.MULADD.X.I.C.16.F	Ensemble multiply add extract immediate signed complex doublets floor
E.MULADD.X.I.C.16.N	Ensemble multiply add extract immediate signed complex doublets nearest
E.MULADD.X.I.C.16.Z	Ensemble multiply add extract immediate signed complex doublets zero
E.MULADD.X.I.C.32.C	Ensemble multiply add extract immediate signed complex quadlets ceiling
E.MULADD.X.I.C.32.F	Ensemble multiply add extract immediate signed complex quadlets floor
E.MULADD.X.I.C.32.N	Ensemble multiply add extract immediate signed complex quadlets nearest
E.MULADD.X.I.C.32.Z	Ensemble multiply add extract immediate signed complex quadlets zero
E.MULADD.X.I.C.64.C	Ensemble multiply add extract immediate signed complex octlets ceiling
E.MULADD.X.I.C.64.F	Ensemble multiply add extract immediate signed complex octlets floor
E.MULADD.X.I.C.64.N	Ensemble multiply add extract immediate signed complex octlets nearest
E.MULADD.X.I.C.64.Z	Ensemble multiply add extract immediate signed complex octlets zero
E.MULADD.X.I.M.8.C	Ensemble multiply add extract immediate mixed-signed bytes ceiling
E.MULADD.X.I.M.8.F	Ensemble multiply add extract immediate mixed-signed bytes floor
E.MULADD.X.I.M.8.N	Ensemble multiply add extract immediate mixed-signed bytes nearest
E.MULADD.X.I.M.8.Z	Ensemble multiply add extract immediate mixed-signed bytes zero
E.MULADD.X.I.M.16.C	Ensemble multiply add extract immediate mixed-signed doublets ceiling
E.MULADD.X.I.M.16.F	Ensemble multiply add extract immediate mixed-signed doublets floor
E.MULADD.X.I.M.16.N	Ensemble multiply add extract immediate mixed-signed doublets nearest
E.MULADD.X.I.M.16.Z	Ensemble multiply add extract immediate mixed-signed doublets zero
E.MULADD.X.I.M.32.C	Ensemble multiply add extract immediate mixed-signed quadlets ceiling
E.MULADD.X.I.M.32.F	Ensemble multiply add extract immediate mixed-signed quadlets floor
E.MULADD.X.I.M.32.N	Ensemble multiply add extract immediate mixed-signed quadlets nearest
E.MULADD.X.I.M.32.Z	Ensemble multiply add extract immediate mixed-signed quadlets zero
E.MULADD.X.I.M.64.C	Ensemble multiply add extract immediate mixed-signed octlets ceiling
E.MULADD.X.I.M.64.F	Ensemble multiply add extract immediate mixed-signed octlets floor
E.MULADD.X.I.M.64.N	Ensemble multiply add extract immediate mixed-signed octlets nearest
E.MULADD.X.I.M.64.Z	Ensemble multiply add extract immediate mixed-signed octlets zero
E.MULADD.X.I.8.C	Ensemble multiply add extract immediate signed bytes ceiling
E.MULADD.X.I.8.F	Ensemble multiply add extract immediate signed bytes floor
E.MULADD.X.I.8.N	Ensemble multiply add extract immediate signed bytes nearest
E.MULADD.X.I.8.Z	Ensemble multiply add extract immediate signed bytes zero
E.MULADD.X.I.16.C	Ensemble multiply add extract immediate signed doublets ceiling
E.MULADD.X.I.16.F	Ensemble multiply add extract immediate signed doublets floor

E.MULADD.X.I.16.N	Ensemble multiply add extract immediate signed doublets nearest
E.MULADD.X.I.16.Z	Ensemble multiply add extract immediate signed doublets zero
E.MULADD.X.I.32.C	Ensemble multiply add extract immediate signed quadlets ceiling
E.MULADD.X.I.32.F	Ensemble multiply add extract immediate signed quadlets floor
E.MULADD.X.I.32.N	Ensemble multiply add extract immediate signed quadlets nearest
E.MULADD.X.I.32.Z	Ensemble multiply add extract immediate signed quadlets zero
E.MULADD.X.I.64.C	Ensemble multiply add extract immediate signed octlets ceiling
E.MULADD.X.I.64.F	Ensemble multiply add extract immediate signed octlets floor
E.MULADD.X.I.64.N	Ensemble multiply add extract immediate signed octlets nearest
E.MULADD.X.I.64.Z	Ensemble multiply add extract immediate signed octlets zero
E.MULADD.X.I.U.8.C	Ensemble multiply add extract immediate unsigned bytes ceiling
E.MULADD.X.I.U.8.F	Ensemble multiply add extract immediate unsigned bytes floor
E.MULADD.X.I.U.8.N	Ensemble multiply add extract immediate unsigned bytes nearest
E.MULADD.X.I.U.16.C	Ensemble multiply add extract immediate unsigned doublets ceiling
E.MULADD.X.I.U.16.F	Ensemble multiply add extract immediate unsigned doublets floor
E.MULADD.X.I.U.16.N	Ensemble multiply add extract immediate unsigned doublets nearest
E.MULADD.X.I.U.32.C	Ensemble multiply add extract immediate unsigned quadlets ceiling
E.MULADD.X.I.U.32.F	Ensemble multiply add extract immediate unsigned quadlets floor
E.MULADD.X.I.U.32.N	Ensemble multiply add extract immediate unsigned quadlets nearest
E.MULADD.X.I.U.64.C	Ensemble multiply add extract immediate unsigned octlets ceiling
E.MULADD.X.I.U.64.F	Ensemble multiply add extract immediate unsigned octlets floor
E.MULADD.X.I.U.64.N	Ensemble multiply add extract immediate unsigned octlets nearest

Format

E.op.size.rnd rd@rc,rb,i

rd=eopsizernd(rd,rc,rb,i)

31	24 23	18 17	12 11	6 5 4 3 2 1 0
E.op	rd	rc	rb	sz rnd sh
8	6	6	6	2 2 2

```

sz ← log(size) - 3
case op of
    E.MULADD.X.I:
        sh ← size - i - 1
    E.MULADD.X.I.U, E.MULADD.X.I.M, E.MULADD.X.I.C:
        sh ← size - i
endcase

```

Description


The contents of registers rc and rb are partitioned into groups of operands of the size specified and multiplied, added or subtracted, or are catenated and partitioned into operands of twice the size specified. The contents of register rd are partitioned into groups of operands of the size specified and sign or zero ensemble and shifted as specified, then added

to the group of values computed. The group of values is rounded, and limited as specified, yielding a group of results which is the size specified. The group of results is catenated and placed in register rd.


For mixed-signed multiplies, the contents of register rc is signed, and the contents of register rb as unsigned. The extraction operation, the contents of register rd, and the result of mixed-signed multiplies are signed.

Z (zero) rounding is not defined for unsigned extract operations, and a ReservedInstruction exception is raised if attempted. F (floor) rounding will properly round unsigned results downward.

An ensemble multiply add extract immediate doublets instruction (E.MUL.ADD.X.I.16 or E.MUL.ADD.X.I.U.16) multiplies operand [h g f e d c b a] by operand [p o n m l k j i], then adding [x w v u t s r q], yielding the products [hp+x go+w fn+v em+u dl+t ck+s bj+r ai+q], rounded and limited as specified:




Another illustration of ensemble multiply add extract immediate doublets instruction (E.MUL.ADDXI.16 or E.MUL.ADD.X.I.U.16):




Ensemble multiply add extract immediate doublets

An ensemble multiply add extract immediate complex doublets instruction (E.MUL.ADD.X.I.C.16 or G.MUL.ADD.X.I.U.16) multiplies operand [h g f e d c b a] by operand [p o n m l k j i], then adding [x w v u t s r q], yielding the result [gp+ho+x go-hp+w en+fm+v em-fn+u cl+dk+t ck-dl+s aj+bi+r ai-bj+q], rounded and limited as specified. Note that this instruction prefers an organization of complex numbers in which the real part is located to the right (lower precision) of the imaginary part.:



Ensemble multiply add extract immediate complex doublets

Another illustration of ensemble multiply extract immediate complex doublets instruction (E.MUL.ADD.X.I.C.16):



Ensemble multiply add extract immediate complex doublets

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
endef

def EnsembleExtractImmediateInplace(op,rnd,size,rd,rc,rb,sh)
    d ← RegRead[rd, 128]
    c ← RegRead[rc, 128]
    b ← RegRead[rb, 128]
    case op of
        E.MULADD.X.I, E.MULADD.X.I.C:
            ds ← 1
            cs ← 1
            bs ← 1
        E.MULADD.X.I.M:
            ds ← 1
            cs ← 0
            bs ← 1
        E.MULADD.X.I.U:
            ds ← 0
            cs ← 0
            bs ← 0
            if rnd = Z then
                raise ReservedInstruction
            endif
    endcase

```

```

case op of
  E.MULADD.X.I, E.MULADD.X.I.U, E.MULADD.X.I.M:
    h ← 2*size + 1
  E.MULADD.X.I.C:
    h ← {2*size} + 2
endcase
r ← h - size - sh - 1 - {cs and bs}
for i ← 0 to 128-size by size
  di ← {(ds and di+size-1)h-size-r || (di+size-1..i)110^r}
  case op of
    E.MULADD.X.I, E.MULADD.X.I.M, E.MULADD.X.I.U:
      p ← mul(size,h,cs,c,i,bs,b,i) + di
    E.MULADD.X.I.C:
      if i & size = 0 then
        p ← mul(size,h,cs,c,i,bs,b,i) - mul(size,h,cs,c,i+size,bs,b,i+size) + di
      else
        p ← mul(size,h,cs,c,i,bs,b,i+size) + mul(size,h,cs,c,i,bs,b,i+size) + di
      endif
    endcase
  case rnd of
    none, N:
      s ← 0h-r || ~pr || prr-1
    Z:
      s ← 0h-r || ph-1
    F:
      s ← 0h
    C:
      s ← 0h-r || 1r
  endcase
  v ← {(ds & ph-1)|| p} + {011s}
  if {vh..r+size = (ds & vr+size-1)h+1-r-size} then
    asize-1+i..i ← vsize-1+r..r
  else
    asize-1+i..i ← ds ? {vh || ~vhsize-1} : 1size
  endif
  endfor
  RegWrite(rd, 128, a)
enddef

```

Exceptions

ReservedInstruction

Ensemble Floating-point

These operations take two values from registers, perform a group of floating-point arithmetic operations on partitions of bits in the operands, and place the catenated results in a register.

Operation codes

E.ADD.F.16	Ensemble add floating-point half
E.ADD.F.16.C	Ensemble add floating-point half ceiling
E.ADD.F.16.F	Ensemble add floating-point half floor
E.ADD.F.16.N	Ensemble add floating-point half nearest
E.ADD.F.16.X	Ensemble add floating-point half exact
E.ADD.F.16.Z	Ensemble add floating-point half zero
E.ADD.F.32	Ensemble add floating-point single
E.ADD.F.32.C	Ensemble add floating-point single ceiling
E.ADD.F.32.F	Ensemble add floating-point single floor
E.ADD.F.32.N	Ensemble add floating-point single nearest
E.ADD.F.32.X	Ensemble add floating-point single exact
E.ADD.F.32.Z	Ensemble add floating-point single zero
E.ADD.F.64	Ensemble add floating-point double
E.ADD.F.64.C	Ensemble add floating-point double ceiling
E.ADD.F.64.F	Ensemble add floating-point double floor
E.ADD.F.64.N	Ensemble add floating-point double nearest
E.ADD.F.64.X	Ensemble add floating-point double exact
E.ADD.F.64.Z	Ensemble add floating-point double zero
E.ADD.F.128	Ensemble add floating-point quad
E.ADD.F.128.C	Ensemble add floating-point quad ceiling
E.ADD.F.128.F	Ensemble add floating-point quad floor
E.ADD.F.128.N	Ensemble add floating-point quad nearest
E.ADD.F.128.X	Ensemble add floating-point quad exact
E.ADD.F.128.Z	Ensemble add floating-point quad zero
E.DIV.F.16	Ensemble divide floating-point half
E.DIV.F.16.C	Ensemble divide floating-point half ceiling
E.DIV.F.16.F	Ensemble divide floating-point half floor
E.DIV.F.16.N	Ensemble divide floating-point half nearest
E.DIV.F.16.X	Ensemble divide floating-point half exact
E.DIV.F.16.Z	Ensemble divide floating-point half zero
E.DIV.F.32	Ensemble divide floating-point single
E.DIV.F.32.C	Ensemble divide floating-point single ceiling
E.DIV.F.32.F	Ensemble divide floating-point single floor
E.DIV.F.32.N	Ensemble divide floating-point single nearest
E.DIV.F.32.X	Ensemble divide floating-point single exact
E.DIV.F.32.Z	Ensemble divide floating-point single zero
E.DIV.F.64	Ensemble divide floating-point double
E.DIV.F.64.C	Ensemble divide floating-point double ceiling

E.DIV.F.64.F	Ensemble divide floating-point double floor
E.DIV.F.64.N	Ensemble divide floating-point double nearest
E.DIV.F.64.X	Ensemble divide floating-point double exact
E.DIV.F.64.Z	Ensemble divide floating-point double zero
E.DIV.F.128	Ensemble divide floating-point quad
E.DIV.F.128.C	Ensemble divide floating-point quad ceiling
E.DIV.F.128.F	Ensemble divide floating-point quad floor
E.DIV.F.128.N	Ensemble divide floating-point quad nearest
E.DIV.F.128.X	Ensemble divide floating-point quad exact
E.DIV.F.128.Z	Ensemble divide floating-point quad zero
E.MUL.C.F.16	Ensemble multiply complex floating-point half
E.MUL.C.F.32	Ensemble multiply complex floating-point single
E.MUL.C.F.64	Ensemble multiply complex floating-point double
E.MUL.F.16	Ensemble multiply floating-point half
E.MUL.F.16.C	Ensemble multiply floating-point half ceiling
E.MUL.F.16.F	Ensemble multiply floating-point half floor
E.MUL.F.16.N	Ensemble multiply floating-point half nearest
E.MUL.F.16.X	Ensemble multiply floating-point half exact
E.MUL.F.16.Z	Ensemble multiply floating-point half zero
E.MUL.F.32	Ensemble multiply floating-point single
E.MUL.F.32.C	Ensemble multiply floating-point single ceiling
E.MUL.F.32.F	Ensemble multiply floating-point single floor
E.MUL.F.32.N	Ensemble multiply floating-point single nearest
E.MUL.F.32.X	Ensemble multiply floating-point single exact
E.MUL.F.32.Z	Ensemble multiply floating-point single zero
E.MUL.F.64	Ensemble multiply floating-point double
E.MUL.F.64.C	Ensemble multiply floating-point double ceiling
E.MUL.F.64.F	Ensemble multiply floating-point double floor
E.MUL.F.64.N	Ensemble multiply floating-point double nearest
E.MUL.F.64.X	Ensemble multiply floating-point double exact
E.MUL.F.64.Z	Ensemble multiply floating-point double zero
E.MUL.F.128	Ensemble multiply floating-point quad
E.MUL.F.128.C	Ensemble multiply floating-point quad ceiling
E.MUL.F.128.F	Ensemble multiply floating-point quad floor
E.MUL.F.128.N	Ensemble multiply floating-point quad nearest
E.MUL.F.128.X	Ensemble multiply floating-point quad exact
E.MUL.F.128.Z	Ensemble multiply floating-point quad zero


Selection

class	op	prec	round/trap
add	EADDF	16 32 64 128	NONE C F N X Z
divide	EDIVF	16 32 64 128	NONE C F N X Z
multiply	EMULF	16 32 64 128	NONE C F N X Z
complex multiply	EMUL.CF	16 32 64	NONE

Format

E.op.prec.round rd=rc,rb

rd=eopprecround[rc,rb]

Description

The contents of registers ra and rb are combined using the specified floating-point operation. The result is placed in register rc. The operation is rounded using the specified rounding option or using round-to-nearest if not specified. If a rounding option is specified, the operation raises a floating-point exception if a floating-point invalid operation, divide by zero, overflow, or underflow occurs, or when specified, if the result is inexact. If a rounding option is not specified, floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754.

Definition

```

def mul(size,v,i,w,j) as
    mul ← fmul(F{size,vsize-1+i..i},F{size,wsize-1+j..ji+prec-1..i}
        bi ← F{prec,bi+prec-1..i}
        case op of
            E.ADD.F:
                ai ← faddr(ci,bi,round)
            E.MUL.F:
                ai ← fmul(ci,bi)
            E.MUL.C.F:
                if (i and prec) then
                    ai ← fadd(mul(prec,c,i,b,i-prec), mul(prec,c,i-prec,b,i)))
                else
                    ai ← fsub(mul(prec,c,i,b,i), mul(prec,c,i+prec,b,i+prec)))
                endif
            E.DIV.F.:
                ai ← fdiv(ci,bi)
        endcase
        ai+prec-1..i ← PackF{prec, ai, round}
    endfor
    RegWrite(rd, 128, a)
enddef

```

Exceptions

Floating-point arithmetic

Ensemble Inplace

These operations take operands from three registers, perform operations on partitions of bits in the operands, and place the concatenated results in the third register.

Operation codes

E.MULADD.8	Ensemble multiply signed bytes add doublets
E.MULADD.16	Ensemble multiply signed doublets add quadlets
E.MULADD.32	Ensemble multiply signed quadlets add octlets
E.MULADD.64	Ensemble multiply signed octlets add hexlet
E.MULADD.C.8	Ensemble multiply complex bytes add doublets
E.MULADD.C.16	Ensemble multiply complex doublets add quadlets
E.MULADD.C.32	Ensemble multiply complex quadlets add octlets
E.MULADD.M.8	Ensemble multiply mixed-signed bytes add doublets
E.MULADD.M.16	Ensemble multiply mixed-signed doublets add quadlets
E.MULADD.M.32	Ensemble multiply mixed-signed quadlets add octlets
E.MULADD.M.64	Ensemble multiply mixed-signed octlets add hexlet
E.MULADD.U.8	Ensemble multiply unsigned bytes add doublets
E.MULADD.U.16	Ensemble multiply unsigned doublets add quadlets
E.MULADD.U.32	Ensemble multiply unsigned quadlets add octlets
E.MULADD.U.64	Ensemble multiply unsigned octlets add hexlet
E.MULSUB.8	Ensemble multiply signed bytes subtract doublets
E.MULSUB.16	Ensemble multiply signed doublets subtract quadlets
E.MULSUB.32	Ensemble multiply signed quadlets subtract octlets
E.MULSUB.64	Ensemble multiply signed octlets subtract hexlet
E.MULSUB.C.8	Ensemble multiply complex bytes subtract doublets
E.MULSUB.C.16	Ensemble multiply complex doublets subtract quadlets
E.MULSUB.C.32	Ensemble multiply complex quadlets subtract octlets
E.MULSUB.M.8	Ensemble multiply mixed-signed bytes subtract doublets
E.MULSUB.M.16	Ensemble multiply mixed-signed doublets subtract quadlets
E.MULSUB.M.32	Ensemble multiply mixed-signed quadlets subtract octlets
E.MULSUB.M.64	Ensemble multiply mixed-signed octlets subtract hexlet
E.MULSUB.U.8	Ensemble multiply unsigned bytes subtract doublets
E.MULSUB.U.16	Ensemble multiply unsigned doublets subtract quadlets
E.MULSUB.U.32	Ensemble multiply unsigned quadlets subtract octlets
E.MULSUB.U.64	Ensemble multiply unsigned octlets subtract hexlet

Selection

class	op	type	prec
multiply	E.MULADD	NONE M U	8 16 32 64
complex multiply	E.MULSUB	C	8 16 32

Format

E.op.size rd=rc,rb

rd=gopsize(rd,rc,rb)

31	24 23	18 17	12 11	6 5	0
E.size	rd	rc	rb	op	
8	6	6	6	6	6

Description

The contents of registers rd, rc and rb are fetched. The specified operation is performed on these operands. The result is placed into register rd.

Register rd is both a source and destination of this instruction.

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def EnsembleInplace(op,size,rd,rc,rb) as
    if size=1 then
        raise ReservedInstruction
    endif
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case op of
        E.MULADD, E.MULSUB, E.MULADD.C, E.MULSUB.C:
            cs ← 1
            bs ← 1
        E.MULADD.M, E.MULSUB.M:
            cs ← 0
            bs ← 1
        E.MULADD.U, E.MULSUB.U:
            cs ← 0
            bs ← 0
    endcase
    h ← 2*size
    for i ← 0 to 64-size by size
        di ← d2*(i+size)-1..2*i
        case op of
            E.MULADD, E.MULADD.U, E.MULADD.M:
                p ← mul(size,h,cs,c,i,bs,b,j) + di
            E.MULADD.C:
                if i & size = 0 then
                    p ← mul(size,h,cs,c,i,bs,b,j) - mul(size,h,cs,c,i+size,bs,b,i+size) + di
                else
                    p ← mul(size,h,cs,c,i,bs,b,i+size) + mul(size,h,cs,c,i,bs,b,i+size) + di
                endif
            E.MULSUB, E.MULSUB.U, E.MULSUB.M:
                p ← mul(size,h,cs,c,i,bs,b,j) - di
    endfor
enddef

```

```
E.MULSUBC:  
    if i & size = 0 then  
        p ← mul(size,h,cs,c,i,bs,b,i) - mul(size,h,cs,c,i+size,bs,b,i+size) - di  
    else  
        p ← mul(size,h,cs,c,i,bs,b,i+size) + mul(size,h,cs,c,i,bs,b,i+size) - di  
    endif  
    endcase  
    a2*(i+size)-1..2*i ← p  
endfor  
RegWrite(rd, 128, a)  
enddef
```

Exceptions

none

Ensemble Inplace Floating-point

These operations take operands from three registers, perform operations on partitions of bits in the operands, and place the concatenated results in the third register.

Operation codes

E.MULADD.C.F.16	Ensemble multiply add complex floating-point half
E.MULADD.C.F.32	Ensemble multiply add complex floating-point single
E.MULADD.C.F.64	Ensemble multiply add complex floating-point double
E.MULADD.F.16	Ensemble multiply add floating-point half
E.MULADD.F.16.C	Ensemble multiply add floating-point half ceiling
E.MULADD.F.16.F	Ensemble multiply add floating-point half floor
E.MULADD.F.16.N	Ensemble multiply add floating-point half nearest
E.MULADD.F.16.X	Ensemble multiply add floating-point half exact
E.MULADD.F.16.Z	Ensemble multiply add floating-point half zero
E.MULADD.F.32	Ensemble multiply add floating-point single
E.MULADD.F.32.C	Ensemble multiply add floating-point single ceiling
E.MULADD.F.32.F	Ensemble multiply add floating-point single floor
E.MULADD.F.32.N	Ensemble multiply add floating-point single nearest
E.MULADD.F.32.X	Ensemble multiply add floating-point single exact
E.MULADD.F.32.Z	Ensemble multiply add floating-point single zero
E.MULADD.F.64	Ensemble multiply add floating-point double
E.MULADD.F.64.C	Ensemble multiply add floating-point double ceiling
E.MULADD.F.64.F	Ensemble multiply add floating-point double floor
E.MULADD.F.64.N	Ensemble multiply add floating-point double nearest
E.MULADD.F.64.X	Ensemble multiply add floating-point double exact
E.MULADD.F.64.Z	Ensemble multiply add floating-point double zero
E.MULADD.F.128	Ensemble multiply add floating-point quad
E.MULADD.F.128.C	Ensemble multiply add floating-point quad ceiling
E.MULADD.F.128.F	Ensemble multiply add floating-point quad floor
E.MULADD.F.128.N	Ensemble multiply add floating-point quad nearest
E.MULADD.F.128.X	Ensemble multiply add floating-point quad exact
E.MULADD.F.128.Z	Ensemble multiply add floating-point quad zero
E.MULSUB.C.F.16	Ensemble multiply subtract complex floating-point half
E.MULSUB.C.F.32	Ensemble multiply subtract complex floating-point single
E.MULSUB.C.F.64	Ensemble multiply subtract complex floating-point double
E.MULSUB.F.16	Ensemble multiply subtract floating-point half
E.MULSUB.F.32	Ensemble multiply subtract floating-point single
E.MULSUB.F.64	Ensemble multiply subtract floating-point double
E.MULSUB.F.128	Ensemble multiply subtract floating-point quad


Selection

class	op	type	prec	round/trap
multiply add	E.MULADD	F	16 32 64 128	NONE C F N X Z
		C.F	16 32 64	NONE
multiply subtract	E.MULSUB	F	16 32 64 128	NONE
		C.F	16 32 64	NONE

Format

E.op.size rd@rc,rb

rd=eopsize(rd,rc,rb)

Description

The contents of registers rd, rc and rb are fetched. The specified operation is performed on these operands. The result is placed into register rd.

Register rd is both a source and destination of this instruction.

Definition

```

def mul(size,v,i,w,j) as
    mul ← fmul(F(size,v|size-1+i..i),F(size,w|size-1+j..j))
endef

def EnsembleInplaceFloatingPoint(op,size,rd,rc,rb) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-size by size
        di ← F(prec,d|+prec-1..i)
        case op of
            E.MULADD.F:
                ai ← fadd(di, mul(prec,c,i,b,i))
            E.MULADD.C.F:
                if {i and prec} then
                    ai ← fadd(di, fadd(mul(prec,c,i,b,i-prec), mul(c,i-prec,b,i)))
                else
                    ai ← fadd(di, fsub(mul(prec,c,i,b,i), mul(prec,c,i+prec,b,i+prec)))
                endif
            E.MULSUB.F:
                ai ← frsub(di, mul(prec,c,i,b,i))
            E.MULSUB.C.F:
                if {i and prec} then
                    ai ← frsub(di, fadd(mul(prec,c,i,b,i-prec), mul(c,i-prec,b,i)))
                else
                    ai ← frsub(di, fsub(mul(prec,c,i,b,i), mul(prec,c,i+prec,b,i+prec)))
                endif
    endfor
endef

```

```
        endif
    endcase
    ai+prec-1..i ← PackF(prec, ai, round)
endfor
RegWrite[rd, 128, a]
enddef
```

Exceptions

none

Ensemble Reversed Floating-point

These operations take two values from registers, perform a group of floating-point arithmetic operations on partitions of bits in the operands, and place the concatenated results in a register.

Operation codes

E.SUB.F.16	Ensemble subtract floating-point half
E.SUB.F.16.C	Ensemble subtract floating-point half ceiling
E.SUB.F.16.F	Ensemble subtract floating-point half floor
E.SUB.F.16.N	Ensemble subtract floating-point half nearest
E.SUB.F.16.Z	Ensemble subtract floating-point half zero
E.SUB.F.16.X	Ensemble subtract floating-point half exact
E.SUB.F.32	Ensemble subtract floating-point single
E.SUB.F.32.C	Ensemble subtract floating-point single ceiling
E.SUB.F.32.F	Ensemble subtract floating-point single floor
E.SUB.F.32.N	Ensemble subtract floating-point single nearest
E.SUB.F.32.Z	Ensemble subtract floating-point single zero
E.SUB.F.32.X	Ensemble subtract floating-point single exact
E.SUB.F.64	Ensemble subtract floating-point double
E.SUB.F.64.C	Ensemble subtract floating-point double ceiling
E.SUB.F.64.F	Ensemble subtract floating-point double floor
E.SUB.F.64.N	Ensemble subtract floating-point double nearest
E.SUB.F.64.Z	Ensemble subtract floating-point double zero
E.SUB.F.64.X	Ensemble subtract floating-point double exact
E.SUB.F.128	Ensemble subtract floating-point quad
E.SUB.F.128.C	Ensemble subtract floating-point quad ceiling
E.SUB.F.128.F	Ensemble subtract floating-point quad floor
E.SUB.F.128.N	Ensemble subtract floating-point quad nearest
E.SUB.F.128.Z	Ensemble subtract floating-point quad zero
E.SUB.F.128.X	Ensemble subtract floating-point quad exact


Selection

class	op	prec	round/trap
set	SET. E LG L GE	16 32 64 128	NONE X
subtract	SUB	16 32 64 128	NONE C F N X Z

Format

E.op.prec.round rd=rb,rc

rd=eopprecround(rb,rc)

Description

The contents of registers rc and rb are combined using the specified floating-point operation. The result is placed in register rd. The operation is rounded using the specified rounding option or using round-to-nearest if not specified. If a rounding option is specified, the operation raises a floating-point exception if a floating-point invalid operation, divide by zero, overflow, or underflow occurs, or when specified, if the result is inexact. If a rounding option is not specified, floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754.

Definition

```
def EnsembleReversedFloatingPoint(op,prec,round,rd,rc,rb) as
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-prec by prec
        ci ← F(prec,ci+prec-1..i)
        bi ← F(prec,bi+prec-1..i)
        ai ← frsubr(ci,-bi, round)
        ai+prec-1..i ← PackF(prec, ai, round)
    endfor
    RegWrite(rd, 128, a)
enddef
```

Exceptions

Floating-point arithmetic

Ensemble Ternary

These operations take three values from registers, perform a group of calculations on partitions of bits of the operands and place the catenated results in a fourth register.


Operation codes

E.MUL.G.8	Ensemble multiply Galois field byte
E.MUL.G.64	Ensemble multiply Galois field octlet

Format

E.MUL.G.size ra=rd,rc,rb

ra=emulgsize(rd,rc,rb)




Description

The contents of registers rd, rc, and rb are fetched. The specified operation is performed on these operands. The result is placed into register ra.

The contents of registers rd and rc are partitioned into groups of operands of the size specified and multiplied in the manner of polynomials. The group of values is reduced modulo the polynomial specified by the contents of register rb, yielding a group of results, each of which is the size specified. The group of results is catenated and placed in register ra.

An ensemble multiply Galois field bytes instruction (E.MULG.8) multiplies operand [d15 d14 d13 d12 d11 d10 d9 d8 d7 d6 d5 d4 d3 d2 d1 d0] by operand [c15 c14 c13 c12 c11 c10 c9 c8 c7 c6 c5 c4 c3 c2 c1 c0], modulo polynomial [q], yielding the results [(d15c15 mod q) (d14c14 mod q) ... (d0c0 mod q):



Ensemble multiply Galois field bytes

Definition

```

def c ← PolyMultiply(size,a,b) as
    p[0] ← 02*size
    for k ← 0 to size-1
        p[k+1] ← p[k] ^ ak ? (0size-k || b || 0k) : 02*size
    endfor
    c ← p[size]
enddef

def c ← PolyResidue(size,a,b) as
    p[0] ← a
    for k ← size-1 to 0 by -1
        p[k+1] ← p[k] ^ p[0]size+k ? (0size-k || 1 || b || 0k) : 02*size
    endfor
    c ← p[size]size-1..0
enddef

def EnsembleTernary(op,size,rd,rc,rb,ra) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    case op of

```

E.MUL.G:

```
    for i ← 0 to 128-size by size
        asize-1+i..i ← PolyResidue(size, PolyMul(size, csize-1+i..i, bsize-1+i..i), dsize-1+i..i)
    endfor
endcase
RegWrite(ra, 128, a)
enddef
```

Exceptions

none

Ensemble Ternary Floating-point

These operations take three values from registers, perform a group of floating-point arithmetic operations on partitions of bits in the operands, and place the concatenated results in a register.

Operation codes

E.SCAL.ADD.F.16	Ensemble scale add floating-point half
E.SCAL.ADD.F.32	Ensemble scale add floating-point single
E.SCAL.ADD.F.64	Ensemble scale add floating-point double


Selection

class	op	prec
scale add	E.SCAL.ADD.F	16 32 64

Format

E.SCALADD.F.size ra=rd,rc,rb

ra=escaladdfsize(rd,rc,rb)



Description

The contents of registers rd and rc are taken to represent a group of floating-point operands. Operands from register rd are multiplied with a floating-point operand taken from the least-significant bits of the contents of register rb and added to operands from register rc multiplied with a floating-point operand taken from the next least-significant bits of the contents of register rb. The results are concatenated and placed in register ra. The results are rounded to the nearest representable floating-point value in a single floating-point operation. Floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754. These instructions cannot select a directed rounding mode or trap on inexact.

Definition

```
def EnsembleFloatingPointTernary(op,prec,rd,rc,rb,ra) as
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 128)
    b ← RegRead(rb, 128)
    for i ← 0 to 128-prec by prec
        di ← F[prec,di+prec-1..i]
        ci ← F[prec,ci+prec-1..i]
        ai ← fadd(fmul(di, F[prec,bprec-1..0]), fmul(ci, F[prec,b2*prec-1..prec]))
        ai+prec-1..i ← PackF(prec, ai, none)
```

```
endfor
RegWrite(ra, 128, a)
enddef
```

Exceptions

none

Ensemble Unary

These operations take operands from a register, perform operations on partitions of bits in the operand, and place the concatenated results in a second register.

Operation codes

E.LOG.MOST.8	Ensemble log of most significant bit signed bytes
E.LOG.MOST.16	Ensemble log of most significant bit signed doublets
E.LOG.MOST.32	Ensemble log of most significant bit signed quadlets
E.LOG.MOST.64	Ensemble log of most significant bit signed octlets
E.LOG.MOST.128	Ensemble log of most significant bit signed hexlet
E.LOG.MOST.U.8	Ensemble log of most significant bit unsigned bytes
E.LOG.MOST.U.16	Ensemble log of most significant bit unsigned doublets
E.LOG.MOST.U.32	Ensemble log of most significant bit unsigned quadlets
E.LOG.MOST.U.64	Ensemble log of most significant bit unsigned octlets
E.LOG.MOST.U.128	Ensemble log of most significant bit unsigned hexlet
E.SUM.8	Ensemble sum signed bytes
E.SUM.16	Ensemble sum signed doublets
E.SUM.32	Ensemble sum signed quadlets
E.SUM.64	Ensemble sum signed octlets
E.SUM.U.1 ²⁶	Ensemble sum unsigned bits
E.SUM.U.8	Ensemble sum unsigned bytes
E.SUM.U.16	Ensemble sum unsigned doublets
E.SUM.U.32	Ensemble sum unsigned quadlets
E.SUM.U.64	Ensemble sum unsigned octlets


Selection

class	op	size
sum	SUM	8 16 32 64
	SUM.U	1 8 16 32 64
log most significant bit	LOG.MOST LOG.MOST.U	8 16 32 64 128

Format

E.op.size rd=rc

rd=eopsiz(rc)



²⁶ E.SUM.U.1 is encoded as E.SUM.U.128.

Description

Values are taken from the contents of register rc. The specified operation is performed, and the result is placed in register rd.

Definition

```

def EnsembleUnary(op,size,rd,rc)
  c ← RegRead(rc, 128)
  case op of
    E.LOG.MOST:
      for i ← 0 to 128-size by size
        if ( $c_{i+size-1..i} = 0$ ) then
           $a_{i+size-1..i} \leftarrow -1$ 
        else
          for j ← 0 to size-1
            if  $c_{size-1+i..j+i} = (c_{size-1+i..j} \parallel \text{not } c_{size-1+i})$  then
               $a_{i+size-1..i} \leftarrow j$ 
            endif
          endfor
        endif
      endfor
    E.LOG.MOSTU:
      for i ← 0 to 128-size by size
        if ( $c_{i+size-1..i} = 0$ ) then
           $a_{i+size-1..i} \leftarrow -1$ 
        else
          for j ← 0 to size-1
            if  $c_{size-1+i..j+i} = (0^{size-1-j} \parallel 1)$  then
               $a_{i+size-1..i} \leftarrow j$ 
            endif
          endfor
        endif
      endfor
    E.SUM:
      p[0] ← 0128
      for i ← 0 to 128-size by size
        p[i+size] ← p[i] + ( $c_{size-1+i..i}^{128-size} \parallel c_{size-1+i..i}$ )
      endfor
      a ← p[128]
    E.SUMU:
      p[0] ← 0128
      for i ← 0 to 128-size by size
        p[i+size] ← p[i] + (0128-size  $\parallel c_{size-1+i..i}$ )
      endfor
      a ← p[128]
  endcase
  RegWrite(rd, 128, a)
enddef

```

Exceptions

none

Ensemble Unary Floating-point

These operations take one value from a register, perform a group of floating-point arithmetic operations on partitions of bits in the operands, and place the concatenated results in a register.

Operation codes

E.ABS.F.16	Ensemble absolute value floating-point half
E.ABS.F.16.X	Ensemble absolute value floating-point half exception
E.ABS.F.32	Ensemble absolute value floating-point single
E.ABS.F.32.X	Ensemble absolute value floating-point single exception
E.ABS.F.64	Ensemble absolute value floating-point double
E.ABS.F.64.X	Ensemble absolute value floating-point double exception
E.ABS.F.128	Ensemble absolute value floating-point quad
E.ABS.F.128.X	Ensemble absolute value floating-point quad exception
E.COPY.F.16	Ensemble copy floating-point half
E.COPY.F.16.X	Ensemble copy floating-point half exception
E.COPY.F.32	Ensemble copy floating-point single
E.COPY.F.32.X	Ensemble copy floating-point single exception
E.COPY.F.64	Ensemble copy floating-point double
E.COPY.F.64.X	Ensemble copy floating-point double exception
E.COPY.F.128	Ensemble copy floating-point quad
E.COPY.F.128.X	Ensemble copy floating-point quad exception
E.DEFLATE.F.32	Ensemble convert floating-point half from single
E.DEFLATE.F.32.C	Ensemble convert floating-point half from single ceiling
E.DEFLATE.F.32.F	Ensemble convert floating-point half from single floor
E.DEFLATE.F.32.N	Ensemble convert floating-point half from single nearest
E.DEFLATE.F.32.X	Ensemble convert floating-point half from single exact
E.DEFLATE.F.32.Z	Ensemble convert floating-point half from single zero
E.DEFLATE.F.64	Ensemble convert floating-point single from double
E.DEFLATE.F.64.C	Ensemble convert floating-point single from double ceiling
E.DEFLATE.F.64.F	Ensemble convert floating-point single from double floor
E.DEFLATE.F.64.N	Ensemble convert floating-point single from double nearest
E.DEFLATE.F.64.X	Ensemble convert floating-point single from double exact
E.DEFLATE.F.64.Z	Ensemble convert floating-point single from double zero
E.DEFLATE.F.128	Ensemble convert floating-point double from quad
E.DEFLATE.F.128.C	Ensemble convert floating-point double from quad ceiling
E.DEFLATE.F.128.F	Ensemble convert floating-point double from quad floor
E.DEFLATE.F.128.N	Ensemble convert floating-point double from quad nearest
E.DEFLATE.F.128.X	Ensemble convert floating-point double from quad exact
E.DEFLATE.F.128.Z	Ensemble convert floating-point double from quad zero
E.FLOAT.F.16	Ensemble convert floating-point half from doublets
E.FLOAT.F.16.C	Ensemble convert floating-point half from doublets ceiling
E.FLOAT.F.16.F	Ensemble convert floating-point half from doublets floor
E.FLOAT.F.16.N	Ensemble convert floating-point half from doublets nearest

E.FLOAT.F.16.X	Ensemble convert floating-point half from doublets exact
E.FLOAT.F.16.Z	Ensemble convert floating-point half from doublets zero
E.FLOAT.F.32	Ensemble convert floating-point single from quadlets
E.FLOAT.F.32.C	Ensemble convert floating-point single from quadlets ceiling
E.FLOAT.F.32.F	Ensemble convert floating-point single from quadlets floor
E.FLOAT.F.32.N	Ensemble convert floating-point single from quadlets nearest
E.FLOAT.F.32.X	Ensemble convert floating-point single from quadlets exact
E.FLOAT.F.32.Z	Ensemble convert floating-point single from quadlets zero
E.FLOAT.F.64	Ensemble convert floating-point double from octlets
E.FLOAT.F.64.C	Ensemble convert floating-point double from octlets ceiling
E.FLOAT.F.64.F	Ensemble convert floating-point double from octlets floor
E.FLOAT.F.64.N	Ensemble convert floating-point double from octlets nearest
E.FLOAT.F.64.X	Ensemble convert floating-point double from octlets exact
E.FLOAT.F.64.Z	Ensemble convert floating-point double from octlets zero
E.FLOAT.F.128	Ensemble convert floating-point quad from hexlet
E.FLOAT.F.128.C	Ensemble convert floating-point quad from hexlet ceiling
E.FLOAT.F.128.F	Ensemble convert floating-point quad from hexlet floor
E.FLOAT.F.128.N	Ensemble convert floating-point quad from hexlet nearest
E.FLOAT.F.128.X	Ensemble convert floating-point quad from hexlet exact
E.FLOAT.F.128.Z	Ensemble convert floating-point quad from hexlet zero
E.INFLATE.F.16	Ensemble convert floating-point single from half
E.INFLATE.F.16.X	Ensemble convert floating-point single from half exception
E.INFLATE.F.32	Ensemble convert floating-point double from single
E.INFLATE.F.32.X	Ensemble convert floating-point double from single exception
E.INFLATE.F.64	Ensemble convert floating-point quad from double
E.INFLATE.F.64.X	Ensemble convert floating-point quad from double exception
E.NEG.F.16	Ensemble negate floating-point half
E.NEG.F.16.X	Ensemble negate floating-point half exception
E.NEG.F.32	Ensemble negate floating-point single
E.NEG.F.32.X	Ensemble negate floating-point single exception
E.NEG.F.64	Ensemble negate floating-point double
E.NEG.F.64.X	Ensemble negate floating-point double exception
E.NEG.F.128	Ensemble negate floating-point quad
E.NEG.F.128.X	Ensemble negate floating-point quad exception
E.RECEST.F.16	Ensemble reciprocal estimate floating-point half
E.RECEST.F.16.X	Ensemble reciprocal estimate floating-point half exception
E.RECEST.F.32	Ensemble reciprocal estimate floating-point single
E.RECEST.F.32.X	Ensemble reciprocal estimate floating-point single exception
E.RECEST.F.64	Ensemble reciprocal estimate floating-point double
E.RECEST.F.64.X	Ensemble reciprocal estimate floating-point double exception
E.RECEST.F.128	Ensemble reciprocal estimate floating-point quad
E.RECEST.F.128.X	Ensemble reciprocal estimate floating-point quad exception
E.RSQREST.F.16	Ensemble floating-point reciprocal square root estimate half
E.RSQREST.F.16.X	Ensemble floating-point reciprocal square root estimate half exact
E.RSQREST.F.32	Ensemble floating-point reciprocal square root estimate single
E.RSQREST.F.32.X	Ensemble floating-point reciprocal square root estimate single exact

E.RSOREST.F.64	Ensemble floating-point reciprocal square root estimate double
E.RSOREST.F.64.X	Ensemble floating-point reciprocal square root estimate double exact
E.RSOREST.F.128	Ensemble floating-point reciprocal square root estimate quad
E.RSOREST.F.128.X	Ensemble floating-point reciprocal square root estimate quad exact
E.SINK.F.16	Ensemble convert floating-point doublets from half nearest default
E.SINK.F.16.C	Ensemble convert floating-point doublets from half ceiling
E.SINK.F.16.C.D	Ensemble convert floating-point doublets from half ceiling default
E.SINK.F.16.F	Ensemble convert floating-point doublets from half floor
E.SINK.F.16.F.D	Ensemble convert floating-point doublets from half floor default
E.SINK.F.16.N	Ensemble convert floating-point doublets from half nearest
E.SINK.F.16.X	Ensemble convert floating-point doublets from half exact
E.SINK.F.16.Z	Ensemble convert floating-point doublets from half zero
E.SINK.F.16.Z.D	Ensemble convert floating-point doublets from half zero default
E.SINK.F.32	Ensemble convert floating-point quadlets from single nearest default
E.SINK.F.32.C	Ensemble convert floating-point quadlets from single ceiling
E.SINK.F.32.C.D	Ensemble convert floating-point quadlets from single ceiling default
E.SINK.F.32.F	Ensemble convert floating-point quadlets from single floor
E.SINK.F.32.F.D	Ensemble convert floating-point quadlets from single floor default
E.SINK.F.32.N	Ensemble convert floating-point quadlets from single nearest
E.SINK.F.32.X	Ensemble convert floating-point quadlets from single exact
E.SINK.F.32.Z	Ensemble convert floating-point quadlets from single zero
E.SINK.F.32.Z.D	Ensemble convert floating-point quadlets from single zero default
E.SINK.F.64	Ensemble convert floating-point octlets from double nearest default
E.SINK.F.64.C	Ensemble convert floating-point octlets from double ceiling
E.SINK.F.64.C.D	Ensemble convert floating-point octlets from double ceiling default
E.SINK.F.64.F	Ensemble convert floating-point octlets from double floor
E.SINK.F.64.F.D	Ensemble convert floating-point octlets from double floor default
E.SINK.F.64.N	Ensemble convert floating-point octlets from double nearest
E.SINK.F.64.X	Ensemble convert floating-point octlets from double exact
E.SINK.F.64.Z	Ensemble convert floating-point octlets from double zero
E.SINK.F.64.Z.D	Ensemble convert floating-point octlets from double zero default
E.SINK.F.128	Ensemble convert floating-point hexlet from quad nearest default
E.SINK.F.128.C	Ensemble convert floating-point hexlet from quad ceiling
E.SINK.F.128.C.D	Ensemble convert floating-point hexlet from quad ceiling default
E.SINK.F.128.F	Ensemble convert floating-point hexlet from quad floor
E.SINK.F.128.F.D	Ensemble convert floating-point hexlet from quad floor default
E.SINK.F.128.N	Ensemble convert floating-point hexlet from quad nearest
E.SINK.F.128.X	Ensemble convert floating-point hexlet from quad exact
E.SINK.F.128.Z	Ensemble convert floating-point hexlet from quad zero
E.SINK.F.128.Z.D	Ensemble convert floating-point hexlet from quad zero default
E.SQR.F.16	Ensemble square root floating-point half
E.SQR.F.16.C	Ensemble square root floating-point half ceiling
E.SQR.F.16.F	Ensemble square root floating-point half floor
E.SQR.F.16.N	Ensemble square root floating-point half nearest
E.SQR.F.16.X	Ensemble square root floating-point half exact
E.SQR.F.16.Z	Ensemble square root floating-point half zero

E.SQR.F.32	Ensemble square root floating-point single
E.SQR.F.32.C	Ensemble square root floating-point single ceiling
E.SQR.F.32.F	Ensemble square root floating-point single floor
E.SQR.F.32.N	Ensemble square root floating-point single nearest
E.SQR.F.32.X	Ensemble square root floating-point single exact
E.SQR.F.32.Z	Ensemble square root floating-point single zero
E.SQR.F.64	Ensemble square root floating-point double
E.SQR.F.64.C	Ensemble square root floating-point double ceiling
E.SQR.F.64.F	Ensemble square root floating-point double floor
E.SQR.F.64.N	Ensemble square root floating-point double nearest
E.SQR.F.64.X	Ensemble square root floating-point double exact
E.SQR.F.64.Z	Ensemble square root floating-point double zero
E.SQR.F.128	Ensemble square root floating-point quad
E.SQR.F.128.C	Ensemble square root floating-point quad ceiling
E.SQR.F.128.F	Ensemble square root floating-point quad floor
E.SQR.F.128.N	Ensemble square root floating-point quad nearest
E.SQR.F.128.X	Ensemble square root floating-point quad exact
E.SQR.F.128.Z	Ensemble square root floating-point quad zero
E.SUM.F.16	Ensemble sum floating-point half
E.SUM.F.16.C	Ensemble sum floating-point half ceiling
E.SUM.F.16.F	Ensemble sum floating-point half floor
E.SUM.F.16.N	Ensemble sum floating-point half nearest
E.SUM.F.16.X	Ensemble sum floating-point half exact
E.SUM.F.16.Z	Ensemble sum floating-point half zero
E.SUM.F.32	Ensemble sum floating-point single
E.SUM.F.32.C	Ensemble sum floating-point single ceiling
E.SUM.F.32.F	Ensemble sum floating-point single floor
E.SUM.F.32.N	Ensemble sum floating-point single nearest
E.SUM.F.32.X	Ensemble sum floating-point single exact
E.SUM.F.32.Z	Ensemble sum floating-point single zero
E.SUM.F.64	Ensemble sum floating-point double
E.SUM.F.64.C	Ensemble sum floating-point double ceiling
E.SUM.F.64.F	Ensemble sum floating-point double floor
E.SUM.F.64.N	Ensemble sum floating-point double nearest
E.SUM.F.64.X	Ensemble sum floating-point double exact
E.SUM.F.64.Z	Ensemble sum floating-point double zero
E.SUM.F.128	Ensemble sum floating-point quad
E.SUM.F.128.C	Ensemble sum floating-point quad ceiling
E.SUM.F.128.F	Ensemble sum floating-point quad floor
E.SUM.F.128.N	Ensemble sum floating-point quad nearest
E.SUM.F.128.X	Ensemble sum floating-point quad exact
E.SUM.F.128.Z	Ensemble sum floating-point quad zero


Selection

	op	prec	round/trap				
copy	COPY	16 32 64 128	NONE	X			
absolute value	ABS	16 32 64 128	NONE	X			
float from integer	FLOAT	16 32 64 128	NONE	C F N X Z			
integer from float	SINK	16 32 64 128	NONE	C F N X Z	C.D	F.D	Z.D
increase format precision	INFLATE	16 32 64	NONE	X			
decrease format precision	DEFLATE	32 64 128	NONE	C F N X Z			
negate	NEG	16 32 64 128	NONE	X			
reciprocal estimate	RECEST	16 32 64 128	NONE	X			
reciprocal square root estimate	RSQREST	16 32 64 128	NONE	X			
square root	SQR	16 32 64 128	NONE	C F N X Z			
sum	SUM	16 32 64 128	NONE	C F N X Z			

Format

E.op.prec.round rd=rc

rd=eopprecround(rc)

Description

The contents of register rc is used as the operand of the specified floating-point operation. The result is placed in register rd.

The operation is rounded using the specified rounding option or using round-to-nearest if not specified. If a rounding option is specified, unless default exception handling is specified, the operation raises a floating-point exception if a floating-point invalid operation, divide by zero, overflow, or underflow occurs, or when specified, if the result is inexact. If a rounding option is not specified or if default exception handling is specified, floating-point exceptions are not raised, and are handled according to the default rules of IEEE 754.

The reciprocal estimate and reciprocal square root estimate instructions compute an exact result for half precision, and a result with at least 12 bits of significant precision for larger formats.

Definition

```

def EnsembleUnaryFloatingPoint(op,prec,round,rd,rc) as
    c ← RegRead(rc, 128)
    case op of
        E.ABS.F, E.NEG.F, E.SQR.F:
            for i ← 0 to 128-prec by prec
                ci ← F[prec,ci+prec-1..i]
                case op of
                    E.ABS.F:
                        ai.t ← ci.t
                        ai.s ← 0
                        ai.e ← ci.e
                        ai.f ← ci.f
                    E.COPY.F:
                        ai ← ci
                    E.NEG.F:
                        ai.t ← ci.t
                        ai.s ← ~ci.s
                        ai.e ← ci.e
                        ai.f ← ci.f
                    E.RECEST.F:
                        ai ← frecest(ci)
                    E.RSQREST.F:
                        ai ← frsqrest(ci)
                    E.SQR.F:
                        ai ← fsqr(ci)
                endcase
                ai+prec-1..i ← PackF(prec, ai, round)
            endfor
        E.SUM.F:
            p[0].t ← NULL
            for i ← 0 to 128-prec by prec
                p[i+prec] ← fadd(p[i], F[prec,ci+prec-1..i])
            endfor
            a ← PackF(prec, p[128], round)
        E.SINK.F:
            for i ← 0 to 128-prec by prec
                ci ← F[prec,ci+prec-1..i]
                ai+prec-1..i ← fsinkr(prec, ci, round)
            endfor
        E.FLOAT.F:
            for i ← 0 to 128-prec by prec
                ci.t ← NORM
                ci.e ← 0
                ci.s ← ci+prec-1
                ci.f ← ci.s ? 1+-ci+prec-2..i : ci+prec-2..i
                ai+prec-1..i ← PackF(prec, ci, round)
            endfor

```

```
E.INFLATE.F:  
    for i ← 0 to 64-prec by prec  
        ci ← F(prec, ci+prec-1..i)  
        ai+prec+prec-1..i ← PackF(prec+prec, ci, round)  
    endfor  
E.DEFLATE.F:  
    for i ← 0 to 128-prec by prec  
        ci ← F(prec, ci+prec-1..i)  
        ai/2+prec/2-1..i/2 ← PackF(prec/2, ci, round)  
    endfor  
    a127..64 ← 0  
endcase  
RegWrite[rd, 128, a]  
enddef
```

Exceptions

Floating-point arithmetic

Wide Multiply Matrix

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.

Operation codes

W.MUL.MAT.8.B	Wide multiply matrix signed byte big-endian
W.MUL.MAT.8.L	Wide multiply matrix signed byte little-endian
W.MUL.MAT.16.B	Wide multiply matrix signed doublet big-endian
W.MUL.MAT.16.L	Wide multiply matrix signed doublet little-endian
W.MUL.MAT.32.B	Wide multiply matrix signed quadlet big-endian
W.MUL.MAT.32.L	Wide multiply matrix signed quadlet little-endian
W.MUL.MAT.C.8.B	Wide multiply matrix signed complex byte big-endian
W.MUL.MAT.C.8.L	Wide multiply matrix signed complex byte little-endian
W.MUL.MAT.C.16.B	Wide multiply matrix signed complex doublet big-endian
W.MUL.MAT.C.16.L	Wide multiply matrix signed complex doublet little-endian
W.MUL.MAT.M.8.B	Wide multiply matrix mixed-signed byte big-endian
W.MUL.MAT.M.8.L	Wide multiply matrix mixed-signed byte little-endian
W.MUL.MAT.M.16.B	Wide multiply matrix mixed-signed doublet big-endian
W.MUL.MAT.M.16.L	Wide multiply matrix mixed-signed doublet little-endian
W.MUL.MAT.M.32.B	Wide multiply matrix mixed-signed quadlet big-endian
W.MUL.MAT.M.32.L	Wide multiply matrix mixed-signed quadlet little-endian
W.MUL.MAT.P.8.B	Wide multiply matrix polynomial byte big-endian
W.MUL.MAT.P.8.L	Wide multiply matrix polynomial byte little-endian
W.MUL.MAT.P.16.B	Wide multiply matrix polynomial doublet big-endian
W.MUL.MAT.P.16.L	Wide multiply matrix polynomial doublet little-endian
W.MUL.MAT.P.32.B	Wide multiply matrix polynomial quadlet big-endian
W.MUL.MAT.P.32.L	Wide multiply matrix polynomial quadlet little-endian
W.MUL.MAT.U.8.B	Wide multiply matrix unsigned byte big-endian
W.MUL.MAT.U.8.L	Wide multiply matrix unsigned byte little-endian
W.MUL.MAT.U.16.B	Wide multiply matrix unsigned doublet big-endian
W.MUL.MAT.U.16.L	Wide multiply matrix unsigned doublet little-endian
W.MUL.MAT.U.32.B	Wide multiply matrix unsigned quadlet big-endian
W.MUL.MAT.U.32.L	Wide multiply matrix unsigned quadlet little-endian

Selection

class	op	type	size	order
multiply	W.MUL.MAT.	NONE M U P	8 16 32	B L
		C	8 16	B L

Format

`W.op.size.order` `rd=rc,rb`

`rd=wopsizeorder(rc,rb)`

31	2423	1817	1211	65	43	0
W.MINOR.order	rd	rc	rb	sz	op	
8	6	6	6	2	4	

Description


The contents of register `rc` is used as a virtual address, and a value of specified size is loaded from memory. A second value is the contents of register `rb`. The values are partitioned into groups of operands of the size specified. The second values are multiplied with the first values, then summed, producing a group of result values. The group of result values is catenated and placed in register `rd`.

The memory-multiply instructions (`W.MUL.MAT`, `W.MUL.MAT.C`, `W.MUL.MAT.M`, `W.MUL.MAT.P`, `W.MUL.MAT.U`) perform a partitioned array multiply of up to 8192 bits, that is 64x128 bits. The width of the array can be limited to 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired size in bytes to the virtual address operand: 4, 2, or 1. The array can be limited vertically to 128, 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired memory operand size in bytes to the virtual address operand.

The virtual address must either be aligned to 1024/gsize bytes (or 512/gsize for `W.MUL.MAT.C`) (with gsize measured in bits), or must be the sum of an aligned address and one-half of the size of the memory operand in bytes and/or one-quarter of the size of the result in bytes. An aligned address must be an exact multiple of the size expressed in bytes. If the address is not valid an “access disallowed by virtual address” exception occurs.


A wide-multiply-octlets instruction (`W.MUL.MAT.type.64`, `type=NONE M U P`) is not implemented and causes a reserved instruction exception, as an ensemble-multiply-sum-octlets instruction (`E.MUL.SUM.type.64`) performs the same operation except that the multiplier is sourced from a 128-bit register rather than memory. Similarly, instead of wide-multiply-complex-quadlets instruction (`W.MUL.MAT.C.32`), one should use an ensemble-multiply-complex-quadlets instruction (`E.MUL.SUM.C.32`).

A wide-multiply-doublets instruction (`W.MUL.MAT`, `W.MUL.MAT.M`, `W.MUL.MAT.P`, `W.MUL.MAT.U`) multiplies memory [$m_{31} m_{30} \dots m_1 m_0$] with vector [h g f e d c b a], yielding products [$hm_{31}+gm_{27}+\dots+bm_7+am_3 \dots hm_{28}+gm_{24}+\dots+bm_4+am_0$]:



Wide multiply matrix

A wide-multiply-matrix-complex-doublets instruction (W.MUL.MAT.C) multiplies memory [m15 m14 ... m1 m0] with vector [h g f e d c b a], yielding products [hm14+gm15+...+bm2+am3 ... hm12+gm13+...+bm0+am1 -hm13+gm12+...-bm1+am0];



Wide multiply matrix complex

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def c ← PolyMultiply(size,a,b) as
    p[0] ← 02*size
    for k ← 0 to size-1
        p[k+1] ← p[k] ^ ak ? (0size-k || b || 0k) : 02*size
    endfor
    c ← p[size]
enddef

def MemoryMultiply(major,op,gsize,rd,rc,rb)
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    lgszie ← log(gsize)
    if C2..lgszie-4 ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    endif
    if C2..lgszie-3 ≠ 0 then
        wszie ← (c and (0-c)) || 04
        t ← c and (c-1)
    else
        wszie ← 64
        t ← a
    endif
    lwszie ← log(wszie)
    if twszie+6-lgszie..lwszie-3 ≠ 0 then
        msize ← (t and (0-t)) || 04
        VirtAddr ← t and (t-1)
    else
        msize ← 128*wszie/gsize
        VirtAddr ← t
    endif
    case major of
        W.MINOR.B:
            order ← B
        W.MINOR.L:
            order ← L
    endcase
    case op of
        W.MUL.MAT.U.8, W.MUL.MAT.U.16, W.MUL.MAT.U.32, W.MUL.MAT.U.64:
            ms ← bs ← 0
        W.MUL.MAT.M.8, W.MUL.MAT.M.16, W.MUL.MAT.M.32, W.MUL.MAT.M.64:
            ms ← 0
            bs ← 1
        W.MUL.MAT.8, W.MUL.MAT.16, W.MUL.MAT.32, W.MUL.MAT.64,
        W.MUL.MAT.C.8, W.MUL.MAT.C.16, W.MUL.MAT.C.32, W.MUL.MAT.C.64:
            ms ← bs ← 1
        W.MUL.MAT.P.8, W.MUL.MAT.P.16, W.MUL.MAT.P.32, W.MUL.MAT.P.64:
    endcase

```

```

m ← LoadMemory(c,VirtAddr,msize,order)
h ← 2*gsize
for i ← 0 to wsize-gsize by gsize
    q[0] ← 02*gsize
    for j ← 0 to vsize-gsize by gsize
        case op of
            W.MUL.MAT.P.8, W.MUL.MAT.P.16, W.MUL.MAT.P.32, W.MUL.MAT.P.64:
                k ← i+wsize*j8..lgszie
                q[j+gsize] ← q[j] ^ PolyMultiply[gsize,mk+gsize-1..k,bj+gsize-1..j]
            W.MUL.MAT.C.8, W.MUL.MAT.C.16, W.MUL.MAT.C.32, W.MUL.MAT.C.64:
                if (~i) & j & gsize = 0 then
                    k ← i-(j&gsize)+wsize*j8..lgszie+1
                    q[j+gsize] ← q[j] + mul[gsize,h,ms,m,k,bs,b,j]
                else
                    k ← i+gsize+wsize*j8..lgszie+1
                    q[j+gsize] ← q[j] - mul[gsize,h,ms,m,k,bs,b,j]
                endif
            W.MUL.MAT.8, W.MUL.MAT.16, W.MUL.MAT.32, W.MUL.MAT.64,
            W.MUL.MAT.M.8, W.MUL.MAT.M.16, W.MUL.MAT.M.32, W.MUL.MAT.M.64,
            W.MUL.MAT.U.8, W.MUL.MAT.U.16, W.MUL.MAT.U.32, W.MUL.MAT.U.64:
                q[j+gsize] ← q[j] + mul[gsize,h,ms,m,i+wsize*j8..lgszie,bs,b,j]
        endfor
        a2*gsize-1+2*i..2*i ← q[vsize]
    endfor
    a127..2*wsize ← 0
    RegWrite[rd, 128, a]
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Wide Multiply Matrix Extract

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.


Operation codes

W.MUL.MAT.X.B	Wide multiply matrix extract big-endian
W.MUL.MAT.X.L	Wide multiply matrix extract little-endian

Format

op ra=rc,rd,rb

ra=op|rc,rd,rb)



Description

The contents of register rc is used as a virtual address, and a value of specified size is loaded from memory. A second value is the contents of register rd. The group size and other parameters are specified from the contents of register rb. The values are partitioned into groups of operands of the size specified and are multiplied and summed, producing a group of values. The group of values is rounded, and limited as specified, yielding a group of results which is the size specified. The group of results is catenated and placed in register ra.

NOTE: The size of this operation is determined from the contents of register rb. The multiplier usage is constant, but the memory operand size is inversely related to the group size. Presumably this can be checked for cache validity.

We also use low order bits of rc to designate a size, which must be consistent with the group size. Because the memory operand is cached, the size can also be cached, thus eliminating the time required to decode the size, whether from rb or from rc.

The wide-multiply-matrix-extract instructions (W.MUL.MAT.X.B, W.MUL.MAT.X.L) perform a partitioned array multiply of up to 16384 bits, that is 128x128 bits. The width of the array can be limited to 128, 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired size in bytes to the virtual address operand: 8, 4, 2, or 1. The array can be limited vertically to 128, 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired memory operand size in bytes to the virtual address operand.

Bits 31..0 of the contents of register rb specifies several parameters which control the manner in which data is extracted. The position and default values of the control fields allows for the source position to be added to a fixed control value for dynamic computation, and allows for the lower 16 bits of the control field to be set for some of the simpler extract cases by a single GCOPYI instruction.

31	2423	16 15 14 13 12 11 10 9 8	0
	fsize	dpos	x s n m l rnd gssp

8 8 1 1 1 1 1 2 9

The table below describes the meaning of each label:

label	bits	meaning
fsize	8	field size
dpos	8	destination position
x	1	reserved
s	1	signed vs. unsigned
n	1	complex vs. real multiplication
m	1	mixed-sign vs. same-sign multiplication
l	1	saturation vs. truncation
rnd	2	rounding
gssp	9	group size and source position

The 9-bit gssp field encodes both the group size, gsize, and source position, spos, according to the formula $gssp = 512 - 4 * gsize + spos$. The group size, gsize, is a power of two in the range 1..128. The source position, spos, is in the range 0..(2*gsize)-1.

The values in the s, n, m, t, and rnd fields have the following meaning:

values	s	n	m	l	rnd
0	unsigned			truncate	F
1	signed	real complex	same-sign mixed-sign	saturate	Z
2					N
3					C


The virtual address must be aligned, that is, it must be an exact multiple of the operand size expressed in bytes. If the address is not aligned an "access disallowed by virtual address" exception occurs.

Z (zero) rounding is not defined for unsigned extract operations, and a ReservedInstruction exception is raised if attempted. F (floor) rounding will properly round unsigned results downward.


A wide-multiply-matrix-extract-doublets instruction (W.MUL.MAT.X.B or W.MUL.MAT.X.L) multiplies memory [m63 m62 m61 ... m2 m1 m0] with vector [h g f e d c b a], yielding the products [am7+bm15+cm23+dm31+em39+fm47+gm55+hm63 ... am2+bm10+cm18+dm26+em34+fm42+gm50+hm58]

am1+bm9+cm17+dm25+em33+fm41+gm49+hm57

am0+bm8+cm16+dm24+em32+fm40+gm48+hm56], rounded and limited as specified:



A wide-multiply-matrix-extract-complex-doublets instruction (W.MUL.MAT.X with n set in rb) multiplies memory [m31 m30 m29 ... m2 m1 m0] with vector [h g f e d c b a], yielding the products [am7+bm6+cm15+dm14+em23+fm22+gm31+hm30 ... am2-bm3+cm10-dm11+em18-fm19+gm26-hm27 am1+bm0+cm9+dm8+em17+fm16+gm25+hm24 am0-bm1+cm8-dm9+em16-f17+gm24-hm25], rounded and limited as specified:



Wide multiply extract matrix complex doublets

Definition

```

def mul(size,h,vs,v,i,ws,w,j) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def WideMultiplyExtractMatrix(op,ra,rb,rc,rd)
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    case b8..0 of
        0..255:
            sgsz ← 128
        256..383:
            sgsz ← 64
        384..447:
            sgsz ← 32
        448..479:
            sgsz ← 16
        480..495:
            sgsz ← 8
        496..503:
            sgsz ← 4

```

```

504..507:
    sgszie ← 2
508..511:
    sgszie ← 1
endcase
l ← b11
m ← b12
n ← b13
signed ← b14
if c3..0 ≠ 0 then
    wsize ← {c and {0..c}} || 04
    t ← c and {c-1}
else
    wsize ← 128
    t ← c
endif
if sgszie < 8 then
    gsize ← 8
elseif sgszie > wsize/2 then
    gsize ← wsize/2
else
    gsize ← sgszie
endif
lgszie ← log(gsize)
lwsize ← log(wsize)
if 4lwsize+6-n-lgszie.lwsize-3 ≠ 0 then
    msize ← {t and {0..t}} || 04
    VirtAddr ← t and {t-1}
else
    msize ← 64*(2-n)*wsize/gsize
    VirtAddr ← t
endif
vsize ← (1+n)*msize*gsize/wsize
mm ← LoadMemory(c,VirtAddr,msize,order)
h ← (2*gsize) + 7 - lgszie
lmsize ← log(msize)
if (VirtAddr,lmsize-4..0 ≠ 0 then
    raise AccessDisallowedByVirtualAddress
endif
case op of
    W.MUL.MAT.X.B:
        order ← B
    W.MUL.MAT.X.L:
        order ← L
endcase
ms ← signed
ds ← signed ^ m
as ← signed or m
spos ← {b8..0} and {2*gsize-1}
dpos ← {0 || b23..16} and {gsize-1}
r ← spos
sfszie ← {0 || b31..24} and {gsize-1}

```

```

tfsiz ← {sfsiz = 0} or ({sfsiz+dpos} > gsize) ? gsize-dpos : sfsiz
fsize ← {tfsiz + spos > h} ? h - spos : tfsiz
if {b10..9 = Z} & ~signed then
    rnd ← F
else
    rnd ← b10..9
endif
for i ← 0 to wsize-gsize by gsize
    q[0] ← 02*gsize+7-lgsize
    for j ← 0 to vsize-gsize by gsize
        if n then
            if {-i} & j & gsize = 0 then
                k ← i-(j&gsize)+wsize*j8..lgsize+1
                q[j+gsize] ← q[j] + mul{gsize,h,ms,mm,k,ds,d,j}
            else
                k ← i+gsize+wsize*j8..lgsize+1
                q[j+gsize] ← q[j] - mul{gsize,h,ms,mm,k,ds,d,j}
            endif
        else
            q[j+gsize] ← q[j] + mul{gsize,h,ms,mm,i+j*wsize/gsize,ds,d,j}
        endif
    endfor
    p ← q[128]
    case rnd of
        none, N:
            s ← 0h-r || ~pr || prr-1
        Z:
            s ← 0h-r || ph-1
        F:
            s ← 0h
        C:
            s ← 0h-r || 1r
    endcase
    v ← {(ds & ph-1) || p} + (011s)
    if {vh..r+fsize = {as & vr+fsize-1}h+1-r-fsize} or not I then
        w ← {as & vr+fsize-1}gsize-fsize-dpos || vfsize-1+r..r || 0dpos
    else
        w ← {s ? {vh || ~vgsize-dpos-1} : 1gsize-dpos} || 0dpos
    endif
    asize-1+i..i ← w
endfor
a127..wsize ← 0
RegWrite(ra, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB

Access detail required by global TB
Local TB miss
Global TB miss

Wide Multiply Matrix Extract Immediate

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.

Operation codes

W.MUL.MAT.X.I.8.C.B	Wide multiply matrix extract immediate signed bytes big-endian ceiling
W.MUL.MAT.X.I.8.C.L	Wide multiply matrix extract immediate signed bytes little-endian ceiling
W.MUL.MAT.X.I.8.F.B	Wide multiply matrix extract immediate signed bytes big-endian floor
W.MUL.MAT.X.I.8.F.L	Wide multiply matrix extract immediate signed bytes little-endian floor
W.MUL.MAT.X.I.8.N.B	Wide multiply matrix extract immediate signed bytes big-endian nearest
W.MUL.MAT.X.I.8.N.L	Wide multiply matrix extract immediate signed bytes little-endian nearest
W.MUL.MAT.X.I.8.Z.B	Wide multiply matrix extract immediate signed bytes big-endian zero
W.MUL.MAT.X.I.8.Z.L	Wide multiply matrix extract immediate signed bytes little-endian zero
W.MUL.MAT.X.I.16.C.B	Wide multiply matrix extract immediate signed doublets big-endian ceiling
W.MUL.MAT.X.I.16.C.L	Wide multiply matrix extract immediate signed doublets little-endian ceiling
W.MUL.MAT.X.I.16.F.B	Wide multiply matrix extract immediate signed doublets big-endian floor
W.MUL.MAT.X.I.16.F.L	Wide multiply matrix extract immediate signed doublets little-endian floor
W.MUL.MAT.X.I.16.N.B	Wide multiply matrix extract immediate signed doublets big-endian nearest
W.MUL.MAT.X.I.16.N.L	Wide multiply matrix extract immediate signed doublets little-endian nearest
W.MUL.MAT.X.I.16.Z.B	Wide multiply matrix extract immediate signed doublets big-endian zero
W.MUL.MAT.X.I.16.Z.L	Wide multiply matrix extract immediate signed doublets little-endian zero
W.MUL.MAT.X.I.32.C.B	Wide multiply matrix extract immediate signed quadlets big-endian ceiling
W.MUL.MAT.X.I.32.C.L	Wide multiply matrix extract immediate signed quadlets little-endian ceiling
W.MUL.MAT.X.I.32.F.B	Wide multiply matrix extract immediate signed quadlets big-endian floor
W.MUL.MAT.X.I.32.F.L	Wide multiply matrix extract immediate signed quadlets little-endian floor
W.MUL.MAT.X.I.32.N.B	Wide multiply matrix extract immediate signed quadlets big-endian nearest
W.MUL.MAT.X.I.32.N.L	Wide multiply matrix extract immediate signed quadlets little-endian nearest
W.MUL.MAT.X.I.32.Z.B	Wide multiply matrix extract immediate signed quadlets big-endian zero
W.MUL.MAT.X.I.32.Z.L	Wide multiply matrix extract immediate signed quadlets little-endian zero
W.MUL.MAT.X.I.64.C.B	Wide multiply matrix extract immediate signed octlets big-endian ceiling
W.MUL.MAT.X.I.64.C.L	Wide multiply matrix extract immediate signed octlets little-endian ceiling
W.MUL.MAT.X.I.64.F.B	Wide multiply matrix extract immediate signed octlets big-endian floor
W.MUL.MAT.X.I.64.F.L	Wide multiply matrix extract immediate signed octlets little-endian floor
W.MUL.MAT.X.I.64.N.B	Wide multiply matrix extract immediate signed octlets big-endian nearest
W.MUL.MAT.X.I.64.N.L	Wide multiply matrix extract immediate signed octlets little-endian nearest
W.MUL.MAT.X.I.64.Z.B	Wide multiply matrix extract immediate signed octlets big-endian zero
W.MUL.MAT.X.I.64.Z.L	Wide multiply matrix extract immediate signed octlets little-endian zero
W.MUL.MAT.X.I.C.8.C.B	Wide multiply matrix extract immediate complex bytes big-endian ceiling
W.MUL.MAT.X.I.C.8.C.L	Wide multiply matrix extract immediate complex bytes little-endian ceiling
W.MUL.MAT.X.I.C.8.F.B	Wide multiply matrix extract immediate complex bytes big-endian floor
W.MUL.MAT.X.I.C.8.F.L	Wide multiply matrix extract immediate complex bytes little-endian floor
W.MUL.MAT.X.I.C.8.N.B	Wide multiply matrix extract immediate complex bytes big-endian nearest
W.MUL.MAT.X.I.C.8.N.L	Wide multiply matrix extract immediate complex bytes little-endian nearest
W.MUL.MAT.X.I.C.8.Z.B	Wide multiply matrix extract immediate complex bytes big-endian zero
W.MUL.MAT.X.I.C.8.Z.L	Wide multiply matrix extract immediate complex bytes little-endian zero

W.MUL.MAT.X.I.C.16.C.B	Wide multiply matrix extract immediate complex doublets big-endian ceiling
W.MUL.MAT.X.I.C.16.C.L	Wide multiply matrix extract immediate complex doublets little-endian ceiling
W.MUL.MAT.X.I.C.16.F.B	Wide multiply matrix extract immediate complex doublets big-endian floor
W.MUL.MAT.X.I.C.16.F.L	Wide multiply matrix extract immediate complex doublets little-endian floor
W.MUL.MAT.X.I.C.16.N.B	Wide multiply matrix extract immediate complex doublets big-endian nearest
W.MUL.MAT.X.I.C.16.N.L	Wide multiply matrix extract immediate complex doublets little-endian nearest
W.MUL.MAT.X.I.C.16.Z.B	Wide multiply matrix extract immediate complex doublets big-endian zero
W.MUL.MAT.X.I.C.16.Z.L	Wide multiply matrix extract immediate complex doublets little-endian zero
W.MUL.MAT.X.I.C.32.C.B	Wide multiply matrix extract immediate complex quadlets big-endian ceiling
W.MUL.MAT.X.I.C.32.C.L	Wide multiply matrix extract immediate complex quadlets little-endian ceiling
W.MUL.MAT.X.I.C.32.F.B	Wide multiply matrix extract immediate complex quadlets big-endian floor
W.MUL.MAT.X.I.C.32.F.L	Wide multiply matrix extract immediate complex quadlets little-endian floor
W.MUL.MAT.X.I.C.32.N.B	Wide multiply matrix extract immediate complex quadlets big-endian nearest
W.MUL.MAT.X.I.C.32.N.L	Wide multiply matrix extract immediate complex quadlets little-endian nearest
W.MUL.MAT.X.I.C.32.Z.B	Wide multiply matrix extract immediate complex quadlets big-endian zero
W.MUL.MAT.X.I.C.32.Z.L	Wide multiply matrix extract immediate complex quadlets little-endian zero
W.MUL.MAT.X.I.C.64.C.B	Wide multiply matrix extract immediate complex octlets big-endian ceiling
W.MUL.MAT.X.I.C.64.C.L	Wide multiply matrix extract immediate complex octlets little-endian ceiling
W.MUL.MAT.X.I.C.64.F.B	Wide multiply matrix extract immediate complex octlets big-endian floor
W.MUL.MAT.X.I.C.64.F.L	Wide multiply matrix extract immediate complex octlets little-endian floor
W.MUL.MAT.X.I.C.64.N.B	Wide multiply matrix extract immediate complex octlets big-endian nearest
W.MUL.MAT.X.I.C.64.N.L	Wide multiply matrix extract immediate complex octlets little-endian nearest
W.MUL.MAT.X.I.C.64.Z.B	Wide multiply matrix extract immediate complex octlets big-endian zero
W.MUL.MAT.X.I.C.64.Z.L	Wide multiply matrix extract immediate complex octlets little-endian zero
W.MUL.MAT.X.I.M.8.C.B	Wide multiply matrix extract immediate mixed-signed bytes big-endian ceiling
W.MUL.MAT.X.I.M.8.C.L	Wide multiply matrix extract immediate mixed-signed bytes little-endian ceiling
W.MUL.MAT.X.I.M.8.F.B	Wide multiply matrix extract immediate mixed-signed bytes big-endian floor
W.MUL.MAT.X.I.M.8.F.L	Wide multiply matrix extract immediate mixed-signed bytes little-endian floor
W.MUL.MAT.X.I.M.8.N.B	Wide multiply matrix extract immediate mixed-signed bytes big-endian nearest
W.MUL.MAT.X.I.M.8.N.L	Wide multiply matrix extract immediate mixed-signed bytes little-endian nearest
W.MUL.MAT.X.I.M.8.Z.B	Wide multiply matrix extract immediate mixed-signed bytes big-endian zero
W.MUL.MAT.X.I.M.8.Z.L	Wide multiply matrix extract immediate mixed-signed bytes little-endian zero
W.MUL.MAT.X.I.M.16.C.B	Wide multiply matrix extract immediate mixed-signed doublets big-endian ceiling
W.MUL.MAT.X.I.M.16.C.L	Wide multiply matrix extract immediate mixed-signed doublets little-endian ceiling
W.MUL.MAT.X.I.M.16.F.B	Wide multiply matrix extract immediate mixed-signed doublets big-endian floor
W.MUL.MAT.X.I.M.16.F.L	Wide multiply matrix extract immediate mixed-signed doublets little-endian floor
W.MUL.MAT.X.I.M.16.N.B	Wide multiply matrix extract immediate mixed-signed doublets big-endian nearest
W.MUL.MAT.X.I.M.16.N.L	Wide multiply matrix extract immediate mixed-signed doublets little-endian nearest
W.MUL.MAT.X.I.M.16.Z.B	Wide multiply matrix extract immediate mixed-signed doublets big-endian zero
W.MUL.MAT.X.I.M.16.Z.L	Wide multiply matrix extract immediate mixed-signed doublets little-endian zero
W.MUL.MAT.X.I.M.32.C.B	Wide multiply matrix extract immediate mixed-signed quadlets big-endian ceiling
W.MUL.MAT.X.I.M.32.C.L	Wide multiply matrix extract immediate mixed-signed quadlets little-endian ceiling
W.MUL.MAT.X.I.M.32.F.B	Wide multiply matrix extract immediate mixed-signed quadlets big-endian floor
W.MUL.MAT.X.I.M.32.F.L	Wide multiply matrix extract immediate mixed-signed quadlets little-endian floor
W.MUL.MAT.X.I.M.32.N.B	Wide multiply matrix extract immediate mixed-signed quadlets big-endian nearest
W.MUL.MAT.X.I.M.32.N.L	Wide multiply matrix extract immediate mixed-signed quadlets little-endian nearest
W.MUL.MAT.X.I.M.32.Z.B	Wide multiply matrix extract immediate mixed-signed quadlets big-endian zero
W.MUL.MAT.X.I.M.32.Z.L	Wide multiply matrix extract immediate mixed-signed quadlets little-endian zero
W.MUL.MAT.X.I.M.64.C.B	Wide multiply matrix extract immediate mixed-signed octlets big-endian ceiling
W.MUL.MAT.X.I.M.64.C.L	Wide multiply matrix extract immediate mixed-signed octlets little-endian ceiling

W.MUL.MAT.X.I.M.64.F.B	Wide multiply matrix extract immediate mixed-signed octlets big-endian floor
W.MUL.MAT.X.I.M.64.F.L	Wide multiply matrix extract immediate mixed-signed octlets little-endian floor
W.MUL.MAT.X.I.M.64.N.B	Wide multiply matrix extract immediate mixed-signed octlets big-endian nearest
W.MUL.MAT.X.I.M.64.N.L	Wide multiply matrix extract immediate mixed-signed octlets little-endian nearest
W.MUL.MAT.X.I.M.64.Z.B	Wide multiply matrix extract immediate mixed-signed octlets big-endian zero
W.MUL.MAT.X.I.M.64.Z.L	Wide multiply matrix extract immediate mixed-signed octlets little-endian zero
W.MUL.MAT.X.I.U.8.C.B	Wide multiply matrix extract immediate unsigned bytes big-endian ceiling
W.MUL.MAT.X.I.U.8.C.L	Wide multiply matrix extract immediate unsigned bytes little-endian ceiling
W.MUL.MAT.X.I.U.8.F.B	Wide multiply matrix extract immediate unsigned bytes big-endian floor
W.MUL.MAT.X.I.U.8.F.L	Wide multiply matrix extract immediate unsigned bytes little-endian floor
W.MUL.MAT.X.I.U.8.N.B	Wide multiply matrix extract immediate unsigned bytes big-endian nearest
W.MUL.MAT.X.I.U.8.N.L	Wide multiply matrix extract immediate unsigned bytes little-endian nearest
W.MUL.MAT.X.I.U.16.C.B	Wide multiply matrix extract immediate unsigned doublets big-endian ceiling
W.MUL.MAT.X.I.U.16.C.L	Wide multiply matrix extract immediate unsigned doublets little-endian ceiling
W.MUL.MAT.X.I.U.16.F.B	Wide multiply matrix extract immediate unsigned doublets big-endian floor
W.MUL.MAT.X.I.U.16.F.L	Wide multiply matrix extract immediate unsigned doublets little-endian floor
W.MUL.MAT.X.I.U.16.N.B	Wide multiply matrix extract immediate unsigned doublets big-endian nearest
W.MUL.MAT.X.I.U.16.N.L	Wide multiply matrix extract immediate unsigned doublets little-endian nearest
W.MUL.MAT.X.I.U.32.C.B	Wide multiply matrix extract immediate unsigned quadlets big-endian ceiling
W.MUL.MAT.X.I.U.32.C.L	Wide multiply matrix extract immediate unsigned quadlets little-endian ceiling
W.MUL.MAT.X.I.U.32.F.B	Wide multiply matrix extract immediate unsigned quadlets big-endian floor
W.MUL.MAT.X.I.U.32.F.L	Wide multiply matrix extract immediate unsigned quadlets little-endian floor
W.MUL.MAT.X.I.U.32.N.B	Wide multiply matrix extract immediate unsigned quadlets big-endian nearest
W.MUL.MAT.X.I.U.32.N.L	Wide multiply matrix extract immediate unsigned quadlets little-endian nearest
W.MUL.MAT.X.I.U.64.C.B	Wide multiply matrix extract immediate unsigned octlets big-endian ceiling
W.MUL.MAT.X.I.U.64.C.L	Wide multiply matrix extract immediate unsigned octlets little-endian ceiling
W.MUL.MAT.X.I.U.64.F.B	Wide multiply matrix extract immediate unsigned octlets big-endian floor
W.MUL.MAT.X.I.U.64.F.L	Wide multiply matrix extract immediate unsigned octlets little-endian floor
W.MUL.MAT.X.I.U.64.N.B	Wide multiply matrix extract immediate unsigned octlets big-endian nearest
W.MUL.MAT.X.I.U.64.N.L	Wide multiply matrix extract immediate unsigned octlets little-endian nearest

Format

W.op.size.rndrd=rc,rb,i

rd=wopsizernd(rc,rb,i)

31	2423	18 17	12 11	6 5	43	21	0
op	rd	rc	rb	sz	rnd	sh	
8	6	6	6	2	2	2	

```

sz ← log(size) - 3
case op of
  W.MUL.MAT.X.I, W.MUL.MAT.X.I.C:
    assert size + 6 - log(size) ≥ i ≥ size + 6 - log(size) - 3
    sh ← size + 6 - log(size) - i
  W.MUL.MAT.X.I.M, W.MUL.MAT.X.I.U:
    assert size + 7 - log(size) ≥ i ≥ size + 7 - log(size) - 3
    sh ← size + 7 - log(size) - i
endcase
  
```

Description


The contents of register rc is used as a virtual address, and a value of specified size is loaded from memory. A second value is the contents of register rb. The values are partitioned into groups of operands of the size specified and are multiplied and summed, or are convolved, producing a group of sums. The group of sums is rounded, and limited as specified, yielding a group of results, each of which is the size specified. The group of results is catenated and placed in register rd.

The wide-multiply-extract-immediate-matrix instructions (W.MUL.MAT.X.I, W.MUL.MAT.X.I.U, W.MUL.MAT.X.I.M, W.MUL.MAT.X.I.C) perform a partitioned array multiply of up to 16384 bits, that is 128x128 bits. The width of the array can be limited to 128, 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired size in bytes to the virtual address operand: 8, 4, 2, or 1. The array can be limited vertically to 128, 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired memory operand size in bytes to the virtual address operand.

The virtual address must either be aligned to 2048/gsize bytes (or 1024/gsize for W.MUL.MAT.X.I.C), or must be the sum of an aligned address and one-half of the size of the memory operand in bytes and/or one-half of the size of the result in bytes. An aligned address must be an exact multiple of the size expressed in bytes. If the address is not valid an “access disallowed by virtual address” exception occurs.


Z (zero) rounding is not defined for unsigned extract operations, and a ReservedInstruction exception is raised if attempted. F (floor) rounding will properly round unsigned results downward.

A wide-multiply-extract-immediate-matrix-doublets instruction (W.MUL.MAT.X.I.16 or W.MUL.MAT.X.I.U.16) multiplies memory [m63 m62 m61 ... m2 m1 m0] with vector [h g f e d c b a], yielding the products [am7+bm15+cm23+dm31+em39+fm47+gm55+hm63 ... am2+bm10+cm18+dm26+em34+fm42+gm50+hm58 am1+bm9+cm17+dm25+em33+fm41+gm49+hm57 am0+bm8+cm16+dm24+em32+fm40+gm48+hm56], rounded and limited as specified:



Wide multiply matrix extract immediate doublets

A wide-multiply-matrix-extract-immediate-complex-doubles instruction (W.MUL.MAT.X.I.C.16) multiplies memory [m31 m30 m29 ... m2 m1 m0] with vector [h g f e d c b a], yielding the products [am7+bm6+cm15+dm14+em23+fm22+gm31+hm30 ... am2-bm3+cm10-dm11+em18-fm19+gm26-hm27 am1+bm0+cm9+dm8+em17+fm16+gm25+hm24 am0-bm1+cm8-dm9+em16-f17+gm24-hm25], rounded and limited as specified:



Wide multiply matrix extract immediate complex doublets

Definition

```

def mul(size,h,vs,v,i,ws,w,jl) as
    mul ← ((vs&vsize-1+i)h-size || vsize-1+i..i) * ((ws&wsize-1+j)h-size || wsize-1+j..j)
enddef

def WideMultiplyExtractImmediateMatrix(op,rnd,gsize,rd,rc,rb,sh)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    lgsiz ← log(gsize)
    case op of
        W.MUL.MAT.X.I.B, WMUL.MAT.X.I.L, WMUL.MAT.X.I.U.B, WMUL.MAT.X.I.U.L,
        WMUL.MAT.X.I.M.B, WMUL.MAT.X.I.M.L :
            if Cgsize-4.0 ≠ 0 then
                raise AccessDisallowedByVirtualAddress
            endif
            if C3..lgsiz-3 ≠ 0 then
                wsize ← (c and (0-c)) || 04
                t ← c and (c-1)
            else
                wsize ← 128
                t ← c
            endif
        else
            if Cgsize-4.0 ≠ 0 then
                raise AccessDisallowedByVirtualAddress
            endif
            if C3..lgsiz-3 ≠ 0 then
                wsize ← (c and (0-c)) || 04
                t ← c and (c-1)
            else
                wsize ← 128
                t ← c
            endif
        endif
    endif
enddef

```

```

        endif
        lwsiz ← log(wsiz)
        if t!wsiz+6-lgsiz..lwsiz-3 ≠ 0 then
            msize ← (t and (0-t)) || 04
            VirtAddr ← t and (t-1)
        else
            msize ← 128*wsiz/gsize
            VirtAddr ← t
        endif
        vsize ← msize*gsize/wsiz
W.MUL.MAT.X.I.C.B, W.MUL.MAT.X.I.C.L:
        if Cgsize-4..0 ≠ 0 then
            raise AccessDisallowedByVirtualAddress
        endif
        if C3..lgsiz-3 ≠ 0 then
            wsiz ← (c and (0-c)) || 04
            t ← c and (c-1)
        else
            wsiz ← 128
            t ← c
        endif
        lwsiz ← log(wsiz)
        if t!wsiz+5-lgsiz..lwsiz-3 ≠ 0 then
            msize ← (t and (0-t)) || 04
            VirtAddr ← t and (t-1)
        else
            msize ← 64*wsiz/gsize
            VirtAddr ← t
        endif
        vsize ← 2*msize*gsize/wsiz
    endcase
    case op of
        W.MUL.MAT.X.I.B, W.MUL.MAT.X.I.U.B, W.MUL.MAT.X.I.M.B, W.MUL.MAT.X.I.C.B:
            order ← B
        W.MUL.MAT.X.I.L, W.MUL.MAT.X.I.U.L, W.MUL.MAT.X.I.M.L, W.MUL.MAT.X.I.C.L:
            order ← L
    endcase
    case op of
        W.MUL.MAT.X.I.U.B, W.MUL.MAT.X.I.U.L:
            as ← ms ← bs ← 0
        W.MUL.MAT.X.I.M.B, W.MUL.MAT.X.I.M.L:
            bs ← 0
            as ← ms ← 1
        W.MUL.MAT.X.I.B, W.MUL.MAT.X.I.L, W.MUL.MAT.X.I.C.B, W.MUL.MAT.X.I.C.L:
            as ← ms ← bs ← 1
    endcase
    m ← LoadMemory(c,VirtAddr,msize,order)
    h ← (2*gsize) + 7 - lgsiz - (ms and bs)
    r ← h - size - sh
    for i ← 0 to wsiz-gsize by gsize
        q[0] ← 02*gsize+7-lgsiz
        for j ← 0 to vsize-gsize by gsize
            case op of

```

```

W.MUL.MAT.X.I.B, W.MUL.MAT.X.I.L,
W.MUL.MAT.X.I.U.B, W.MUL.MAT.X.I.U.L,
W.MUL.MAT.X.I.M.B, W.MUL.MAT.X.I.M.L:
    q[j+gsize] ← q[j] + mul(gsize,h,ms,m,i+wsize*j8..lgsize,bs,b,j)
W.MUL.MAT.X.I.C.B, W.MUL.MAT.X.I.C.L:
    if (-i) & j & gsize = 0 then
        k ← i-(j&gsize)+wsize*j8..lgsize+1
        q[j+gsize] ← q[j] + mul(gsize,h,ms,m,k,bs,b,j)
    else
        k ← i+gsize+wsize*j8..lgsize+1
        q[j+gsize] ← q[j] - mul(gsize,h,ms,m,k,bs,b,j)
    endif
endcase
endfor
p ← q[vsize]
case rnd of
    none, N:
        s ← 0h-r || -pr || prr-1
    Z:
        s ← 0h-r || ph-1
    F:
        .
        s ← 0h
    C:
        s ← 0h-r || 1r
endcase
v ← ((as & ph-1) || p) + (011s)
if (vh..r+gsize = (as & vr+gsize-1)h+1-r-gsize then
    agsize-1+i..i ← vgsize-1+r..r
else
    agsize-1+i..i ← as ? (vh || -vgsize-1) : 1gsize
endif
endfor
a127..wsize ← 0
RegWrite(rd, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Wide Multiply Matrix Floating-point

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.


Operation codes

W.MUL.MAT.C.F.16.B	Wide multiply matrix complex floating-point half big-endian
W.MUL.MAT.C.F.16.L	Wide multiply matrix complex floating-point half little-endian
W.MUL.MAT.C.F.32.B	Wide multiply matrix complex floating-point single big-endian
W.MUL.MAT.C.F.32.L	Wide multiply matrix complex floating-point single little-endian
W.MUL.MAT.C.F.64.B	Wide multiply matrix complex floating-point double big-endian
W.MUL.MAT.C.F.64.L	Wide multiply matrix complex floating-point double little-endian
W.MUL.MAT.F.16.B	Wide multiply matrix floating-point half big-endian
W.MUL.MAT.F.16.L	Wide multiply matrix floating-point half little-endian
W.MUL.MAT.F.32.B	Wide multiply matrix floating-point single big-endian
W.MUL.MAT.F.32.L	Wide multiply matrix floating-point single little-endian
W.MUL.MAT.F.64.B	Wide multiply matrix floating-point double big-endian
W.MUL.MAT.F.64.L	Wide multiply matrix floating-point double little-endian

Format

M.op.size.order rd=rc,rb

rd=mopsizeorder(rc,rb)




Description

The contents of register rc is used as a virtual address, and a value of specified size is loaded from memory. A second value is the contents of register rb. The values are partitioned into groups of operands of the size specified. The second values are multiplied with the first values, then summed, producing a group of result values. The group of result values is catenated and placed in register rd.

The wide-multiply-matrix-floating-point instructions (W.MUL.MAT.F, W.MUL.MAT.C.F) perform a partitioned array multiply of up to 16384 bits, that is 128x128 bits. The width of the array can be limited to 128, 64, 32 bits, but not smaller than twice the group size, by adding one-half the desired size in bytes to the virtual address operand: 8, 4, or 2. The array can be limited vertically to 128, 64, 32, or 16 bits, but not smaller than twice the group size, by adding one-half the desired memory operand size in bytes to the virtual address operand.


The virtual address must either be aligned to 2048/gsize bytes (or 1024/gsize for W.MUL.MAT.C.F), or must be the sum of an aligned address and one-half of the size of the memory operand in bytes and/or one-half of the size of the result in bytes. An aligned address must be an exact multiple of the size expressed in bytes. If the address is not valid an "access disallowed by virtual address" exception occurs.

A wide-multiply-matrix-floating-point-half instruction (W.MUL.MAT.F) multiplies memory [m31 m30 ... m1 m0] with vector [h g f e d c b a], yielding products [hm31+gm27+...+bm7+am3 ... hm28+gm24+...+bm4+am0]:



Wide multiply matrix floating-point half

A wide-multiply-matrix-complex-floating-point-half instruction (W.MUL.MAT.F) multiplies memory [m15 m14 ... m1 m0] with vector [h g f e d c b a], yielding products [hm14+gm15+...+bm2+am3 ... hm12+gm13+...+bm0+am1 -hm13+gm12+...-bm1+am0]:



Wide multiply matrix complex floating-point half

Definition

```

def mul(size,v,i,w,j) as
    mul ← fmul[F[size,vsize-1+i..i],F[size,wsize-1+j..j]]
endef

def MemoryFloatingPointMultiply(major,op,gsize,rd,rc,rb)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    lgsiz ← log(gsize)
    switch op of
        W.MUL.MAT.F.16, W.MUL.MAT.F.32, W.MUL.MAT.F.64:
            if c.lgsiz-4..0 ≠ 0 then
                raise AccessDisallowedByVirtualAddress
            endif
            if c3..lgsiz-3 ≠ 0 then
                wsize ← (c and (0-c)) || 04
                t ← c and (c-1)
            else
                wsize ← 128
                t ← c
            endif
            lwsiz ← log(wsize)
            if t.lwsiz+6-lgsiz..lwsiz-3 ≠ 0 then
                msize ← (t and (0-t)) || 04
                VirtAddr ← t and (t-1)
            else
                msize ← 128*wsize/gsize
                VirtAddr ← t
            endif
    endif
endef

```

```

vsize ← msize*gsize/wsize
W.MUL.MAT.C.F.16, W.MUL.MAT.C.F.32, W.MUL.MAT.C.F.64:
  if Clgsize-4..0 ≠ 0 then
    raise AccessDisallowedByVirtualAddress
  endif
  if C3..lgsiz-3 ≠ 0 then
    wsize ← (c and (0-c)) || 04
    t ← c and (c-1)
  else
    wsize ← 128
    t ← c
  endif
  lwsiz ← log(wsize)
  if lwsiz+5-lgsiz..lwsiz-3 ≠ 0 then
    msize ← (t and (0-t)) || 04
    VirtAddr ← t and (t-1)
  else
    msize ← 64*wsize/gsize
    VirtAddr ← t
  endif
  vsize ← 2*msize*gsize/wsize
endcase
case major of
  M.MINOR.B:
    order ← B
  M.MINOR.L:
    order ← L
endcase
m ← LoadMemory(c,VirtAddr,msize,order)
for i ← 0 to wsize-gsize by gsize
  q[0].t ← NULL
  for j ← 0 to vsize-gsize by gsize
    case op of
      W.MUL.MAT.F.16, W.MUL.MAT.F.32, W.MUL.MAT.F.64:
        q[j+gsize] ← fadd(q[j], mul(gsize,m,i+wsize*j8..lgsiz,b,j))
      W.MUL.MAT.C.F.16, W.MUL.MAT.C.F.32, M.MUL.MAT.C.F.64:
        if (~i) & j & gsize = 0 then
          k ← i-(j&gsize)+wsize*j8..lgsiz+1
          q[j+gsize] ← fadd(q[j], mul(gsize,m,k,b,j))
        else
          k ← i+gsize+wsize*j8..lgsiz+1
          q[j+gsize] ← fsub(q[j], mul(gsize,m,k,b,j))
        endif
    endcase
  endfor
  a[gsize-1+i..i] ← q[vsize]
endfor
a127..wsize ← 0
RegWrite(rd, 128, a)
enddef

```

Exceptions

Floating-point arithmetic
Access disallowed by virtual address
Access disallowed by tag
Access disallowed by global TB
Access disallowed by local TB
Access detail required by tag
Access detail required by local TB
Access detail required by global TB
Local TB miss
Global TB miss

Wide Multiply Matrix Galois

These instructions take an address from a general register to fetch a large operand from memory, second and third operands from general registers, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.

Operation codes

W.MUL.MAT.G.B	Wide multiply matrix Galois big-endian
W.MUL.MAT.G.L	Wide multiply matrix Galois little-endian

Format

W.MUL.MAT.G.order $ra = rc, rd, rb$

$ra = mgmorder(rc, rd, rb)$

31	2423	18 17	12 11	6 5	0
W.MULG.order	rd	rc	rb	ra	
8	6	6	6	6	6


Description

The contents of register rc is used as a virtual address, and a value of specified size is loaded from memory. Second and third values are the contents of registers rd and rb. The values are partitioned into groups of operands of the size specified. The second values are multiplied as polynomials with the first value, producing a result which is reduced to the Galois field specified by the third value, producing a group of result values. The group of result values is catenated and placed in register ra.

The wide-multiply-matrix-Galois instruction (W.MUL.MAT.G) performs a partitioned array multiply of up to 16384 bits, that is 128x128 bits. The width of the array can be limited to 128, 64, 32, or 16 bits, but not smaller than twice the group size of 8 bits, by adding one-half the desired size in bytes to the virtual address operand: 8, 4, 2, or 1. The array can be limited vertically to 128, 64, 32, or 16 bits, but not smaller than twice the group size of 8 bits, by adding one-half the desired memory operand size in bytes to the virtual address operand.

The virtual address must either be aligned to 256 bytes, or must be the sum of an aligned address and one-half of the size of the memory operand in bytes and/or one-half of the size of the result in bytes. An aligned address must be an exact multiple of the size expressed in bytes. If the address is not valid an "access disallowed by virtual address" exception occurs.

A wide-multiply-matrix-Galois instruction (W.MUL.MAT.G) multiplies memory [m255 m254 ... m1 m0] with vector [p o n m l k j i h g f e d c b a], reducing the result modulo polynomial [q], yielding products [(pm255+om247+...+bm31+am15 mod q) (pm254+om246+...+bm30+am14 mod q) ... (pm248+om240+...+bm16+am0 mod q)]:



Wide multiply matrix Galois

Definition

```

def c ← PolyMultiply(size,a,b) as
    p[0] ← 02*size
    for k ← 0 to size-1
        p[k+1] ← p[k] ^ ak ? (0size-k || b || 0k) : 02*size
    endfor
    c ← p[size]
enddef

def c ← PolyResidue(size,a,b) as
    p[0] ← a
    for k ← size-1 to 0 by -1
        p[k+1] ← p[k] ^ p[0]size+k ? (0size-k || 11 || b || 0k) : 02*size
    endfor
    c ← p[size]size-1..0
enddef

def WideMultiplyGalois(op,rd,rc,rb,ra)
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    gsize ← 8

```

```

lysize <- log(lysize)
if Csize-4..0 ≠ 0 then
  raise AccessDisallowedByVirtualAddress
endif
if C3.lysize-3 ≠ 0 then
  wsize <- {c and (0-c)} || 04
  t <- c and (c-1)
else
  wsize <- 128
  t <- c
endif
lysize <- log(wsize)
if tlysize+6-lysize..lysize-3 ≠ 0 then
  msize <- {t and (0-t)} || 04
  VirtAddr <- t and (t-1)
else
  msize <- 128*wsize/gsize
  VirtAddr <- t
endif
case op of
  W.MUL.MAT.G.B:
    order <- B
  W.MUL.MAT.G.L:
    order <- L
endcase
m <- LoadMemory(c,VirtAddr,msize,order)
for i <- 0 to wsize-gsize by gsize
  q[0] <- 02*gsize
  for j <- 0 to vsize-gsize by gsize
    k <- i+wsize*j8..lysize
    q[j+gsize] <- q[j] ^ PolyMultiply(gsize,m[k+gsize-1..k,dj+gsize-1..j])
  endfor
  a[gsize-1+i..i] <- PolyResidue[gsize,q[vsize],b,gsize-1..0]
endfor
a127..wsize <- 0
RegWrite(ra, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Wide Switch

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.


Operation codes

W.SWITCH.B	Wide switch big-endian
W.SWITCH.L	Wide switch little-endian

Format

op ra=rc,rd,rb

ra=op(rc,rd,rb)



Description

The contents of register rc is specified as a virtual address and optionally an operand size, and a value of specified size is loaded from memory. A second value is the catenated contents of registers rd and rb. Eight corresponding bits from the memory value are used to select a single result bit from the second value, for each corresponding bit position. The group of results is catenated and placed in register ra.

The virtual address must either be aligned to 128 bytes, or must be the sum of an aligned address and one-half of the size of the memory operand in bytes. An aligned address must be an exact multiple of the size expressed in bytes. The size of the memory operand must be 8, 16, 32, 64, or 128 bytes. If the address is not valid an “access disallowed by virtual address” exception occurs. When a size smaller than 128 bits is specified, the high order bits of the memory operand are replaced with values corresponding to the bit position, so that the same memory operand specifies a bit selection within symbols of the operand size, and the same operation is performed on each symbol.

Definition

```
def WideSwitch(op,rd,rc,rb,ra)
    d ← RegRead(rd, 128)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    if c1..0 ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    elseif c6..0 ≠ 0 then
        VirtAddr ← c and (c-1)
```

```

w ← wsize ← (c and (0-c)) || 01
else
    VirtAddr ← c
    w ← wsize ← 128
endif
msize ← 8*wsize
lwszie ← log(wsize)
case op of
    W.SWITCH.B:
        order ← B
    W.SWITCH.L:
        order ← L
endcase
m ← LoadMemory(c,VirtAddr,msize,order)
db ← d || b
for i ← 0 to 127
    j ← 0 || i_lwszie-1..0
    k ← m7*w+j || m6*w+j || m5*w+j || m4*w+j || m3*w+j || m2*w+j || m1*w+j || m0
    l ← i_lwszie || j_lwszie-1..0
    aj ← dbl
endfor
RegWrite(ra, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Wide Translate

These instructions take an address from a general register to fetch a large operand from memory, a second operand from a general register, perform a group of operations on partitions of bits in the operands, and catenate the results together, placing the result in a general register.

Operation codes

W.TRANSLATE.8.B	Wide translate bytes big-endian
W.TRANSLATE.16.B	Wide translate doublets big-endian
W.TRANSLATE.32.B	Wide translate quadlets big-endian
W.TRANSLATE.64.B	Wide translate octlets big-endian
W.TRANSLATE.8.L	Wide translate bytes little-endian
W.TRANSLATE.16.L	Wide translate doublets little-endian
W.TRANSLATE.32.L	Wide translate quadlets little-endian
W.TRANSLATE.64.L	Wide translate octlets little-endian

Format

W.TRANSLATE.size.order rd=rc,rb

rd=wtranslate size order(rc,rb)

31	2423	1817	1211	65 4 3	0
W.TRANSLATE.order	rd	rc	rb	sz	o

8

6

6

6

2

4

Description

The contents of register rc is used as a virtual address, and a value of specified size is loaded from memory. A second value is the contents of register rb. The values are partitioned into groups of operands of a size specified. The low-order bytes of the second group of values are used as addresses to choose entries from one or more tables constructed from the first value, producing a group of values. The group of results is catenated and placed in register rd.

By default, the total width of tables is 128 bits, and a total table width of 128, 64, 32, 16 or 8 bits, but not less than the group size may be specified by adding the desired total table width in bytes to the specified address: 16, 8, 4, 2, or 1. When fewer than 128 bits are specified, the tables repeat to fill the 128 bit width.

The default depth of each table is 256 entries, or in bytes is 32 times the group size in bits. An operation may specify 4, 8, 16, 32, 64, 128 or 256 entry tables, by adding one-half of the memory operand size to the address. Table index values are masked to ensure that only the specified portion of the table is used. Tables with just 2 entries cannot be specified; if 2-entry tables are desired, it is recommended to load the entries into registers and use G.MUX to select the table entries.

Failing to initialize the entire table is a potential security hole, as an instruction in with a small-depth table could access table entries previously initialized by an instruction with a large-depth table. We could close this hole either by initializing the entire table, even if extra cycles are required, or by masking the index bits so that only the initialized portion of the table is used. Initializing the entire table with no penalty in cycles could require writing to as many as 128 entries at once, which is quite likely to cause circuit complications. Initializing the entire table with writes to only one entry at a time requires writing 256 cycles, even when the table is smaller. Masking the index bits is the preferred solution.

Masking the index bits suggests that this instruction, for tables larger than 256 entries, may be useful for a general-purpose memory translate function where the processor performs enough independent load operations to fill the 128 bits. Thus, the 16, 32, and 64 bit versions of this function perform equivalent of 8, 4, 2 withdraw, 8, 4, or 2 load-indexed and 7, 3, or 1 group-extract instructions. In other words, this instruction can be as powerful as 23, 11, or 5 existing instructions. The 8-bit version is a single-cycle operation replacing 47 existing instructions, so these are not as big a win, but nonetheless, this is at least a 50% improvement on a 2-issue processor, even with one-cycle-per load timing. To make this possible, the default table size would become 65536, 2³² and 2⁶⁴ for 16, 32 and 64-bit versions of the instruction.

For the big-endian version of this instruction, in the definition below, the contents of register rb is complemented. This reflects a desire to organize the table so that the lowest addressed table entries are selected when the index is zero. In the logical implementation, complementing the index can be avoided by loading the table memory differently for big-endian and little-endian versions. A consequence of this shortcut is that a table loaded by a big-endian translate instruction cannot be used by a little-endian translate instruction, and vice-versa.

The virtual address must either be aligned to 4096 bytes, or must be the sum of an aligned address and one-half of the size of the memory operand in bytes and/or the desired total table width in bytes. An aligned address must be an exact multiple of the size expressed in bytes. The size of the memory operand must be a power of two from 4 to 4096 bytes, but must be at least 4 times the group size and 4 times the total table width. If the address is not valid an “access disallowed by virtual address” exception occurs.

Definition

```
def WideTranslate(op, gsize, rd, rc, rb)
    c ← RegRead(rc, 64)
    b ← RegRead(rb, 128)
    lgsiz ← log(gsize)
    if c.lgsize-4..0 ≠ 0 then
        raise AccessDisallowedByVirtualAddress
    endif
    if c4..lgsiz-3 ≠ 0 then
        wsize ← (c and (0-c)) 11 03
        t ← c and (c-1)
    else
```

```

wsiz ← 128
t ← c
endif
lwsiz ← log(wsiz)
if tlwsiz+4..lwsiz-2 ≠ 0 then
    msize ← (t and {0-t}) || 04
    VirtAddr ← t and {t-1}
else
    msize ← 256*wsiz
    VirtAddr ← t
endif
case op of
    W.TRANSLATE.B:
        order ← B
    W.TRANSLATE.L:
        order ← L
endcase
m ← LoadMemory(c,VirtAddr,msize,order)
vsize ← msize/wsiz
lvsiz ← log(vsize)
for i ← 0 to 128-gsize by gsize
    j ← [(order=B)lvsiz]^(b[lvsiz-1+i..i])*wsiz+i[lwsiz-1..0
    a[gsize-1+i..i] ← mj+gsize-1..j
endfor
RegWrite(rd, 128, a)
enddef

```

Exceptions

Access disallowed by virtual address
 Access disallowed by tag
 Access disallowed by global TB
 Access disallowed by local TB
 Access detail required by tag
 Access detail required by local TB
 Access detail required by global TB
 Local TB miss
 Global TB miss

Memory Management

This section discusses the caches, the translation mechanisms, the memory interfaces, and how the multiprocessor interface is used to maintain cache coherence.

Overview


The Zeus processor provides for both local and global virtual addressing, arbitrary page sizes, and coherent-cache multiprocessing. The memory management system is designed to provide the requirements for implementation of virtual machines as well as virtual memory.

All facilities of the memory management system are themselves memory mapped, in order to provide for the manipulation of these facilities by high-level language, compiled code.

The translation mechanism is designed to allow full byte-at-a-time control of access to the virtual address space, with the assistance of fast exception handlers.

Privilege levels provide for the secure transition between insecure user code and secure system facilities. Instructions execute at a privilege, specified by a two-bit field in the access information. Zero is the least-privileged level, and three is the most-privileged level.

The diagram below sketches the basic organization of the memory management system:



In general terms, the memory management starts from a local virtual address. The local virtual address is translated to a global virtual address by a LTB (Local Translation Buffer). In turn, the global virtual address is translated to a physical address by a GTB (Global

Translation Buffer). One of the addresses, a local virtual address, a global virtual address, or a physical address, is used to index the cache data and cache tag arrays, and one of the addresses is used to check the cache tag array for cache presence. Protection information is assembled from the LTB, GTB, and optionally the cache tag, to determine if the access is legal.

This form varies somewhat, depending on implementation choices made. Because the LTB leaves the lower 48 bits of the address alone, indexing of the cache arrays with the local virtual address is usually identical to cache arrays indexed by the global virtual address. However, indexing cache arrays by the global virtual address rather than the physical address produces a coherence issue if the mapping from global virtual address to physical is many-to-one.

Starting from a local virtual address, the memory management system performs three actions in parallel: the low-order bits of the virtual address are used to directly access the data in the cache, a low-order bit field is used to access the cache tag, and the high-order bits of the virtual address are translated from a local address space to a global virtual address space.

Following these three actions, operations vary depending upon the cache implementation. The cache tag may contain either a physical address and access control information (a physically-tagged cache), or may contain a global virtual address and global protection information (a virtually-tagged cache).

For a physically-tagged cache, the global virtual address is translated to a physical address by the GTB, which generates global protection information. The cache tag is checked against the physical address, to determine a cache hit. In parallel, the local and global protection information is checked.

For a virtually-tagged cache, the cache tag is checked against the global virtual address, to determine a cache hit, and the local and global protection information is checked. If the cache misses, the global virtual address is translated to a physical address by the GTB, which also generates the global protection information.

Local Translation Buffer

The 64-bit global virtual address space is global among all tasks. In a multitask environment, requirements for a task-local address space arise from operations such as the UNIX "fork" function, in which a task is duplicated into parent and child tasks, each now having a unique virtual address space. In addition, when switching tasks, access to one task's address space must be disabled and another task's access enabled.


Zeus provides for portions of the address space to be made local to individual tasks, with a translation to the global virtual space specified by four 16-bit registers for each local virtual space. The registers specify a mask selecting which of the high-order 16 address bits are checked to match a particular value, and if they match, a value with which to modify the virtual address. Zeus avoids setting a fixed page size or local address size; these can be set by software conventions.

A local virtual address space is specified by the following:

field name	size	description
lm	16	mask to select fields of local virtual address to perform match over
la	16	value to perform match with masked local virtual address
lx	16	value to xor with local virtual address if matched
lp	16	local protection field (detailed later)
local virtual address space specifiers		

Physical address

There are as many LTB as threads, and up to 2^3 (8) entries per LTB. Each entry is 128 bits, with the high order 64 bits reserved. The physical address of a LTB entry for thread th, entry en, byte b is:



Definition


```

def data,flags ← AccessPhysicalLTB(pa,op,wdata) as
    th ← pa23..19
    en ← pa6..4
    if {en < (1 11 0LE)} and {th < T} and {pa18..6=0} then
        case op of
            R:
                data ← 064 || LTBArray[th][en]
            W:
                LocalTB[th][en] ← wdata63..0
        endcase
    else
        data ← 0
    endif
enddef

```

Entry Format

These 16-bit values are packed together into a 64-bit LTB entry as follows:




The LTB contains a separate context of register sets for each thread, indicated by the th index above. A context consists of one or more sets of lm/la/lx/lp registers, one set for each simultaneously accessible local virtual address range, indicated by the en index above. This set of registers is called the "Local TB context," or LTB (Local Translation Buffer)

context. The effect of this mechanism is to provide the facilities normally attributed to segmentation. However, in this system there is no extension of the address range, instead, segments are local nicknames for portions of the global virtual address space.

A failure to match a LTB entry results either in an exception or an access to the global virtual address space, depending on privilege level. A single bit, selected by the privilege level active for the access from a four bit **control register** field, **global access**, ga determines the result. If **gaPL** is zero (0), the failure causes an exception, if it is one (1), the failure causes the address to be directly used as a global virtual address without modification.

Global Access (fields of control register)



Usually, global access is a right conferred to highly privilege levels, so a typical system may be configured with **ga0** and **ga1** clear (0), but **ga2** and **ga3** set (1). A single low-privilege (0) task can be safely permitted to have global access, as accesses are further limited by the **rwxg** privilege fields. A concrete example of this is an emulation task, which may use global addresses to simulate segmentation, such as an x86 emulation. The emulation task then runs as privilege 0, with **ga0** set, while most user tasks run as privilege 1, with **ga1** clear. Operating system tasks then use privilege 2 and 3 to communicate with and control the user tasks, with **ga2** and **ga3** set.

For tasks that have global access disabled at their current privilege level, failure to match a LTB entry causes an exception. The exception handler may load an LTB entry and continue execution, thus providing access to an arbitrary number of local virtual address ranges.


When failure to match a LTB entry does not cause an exception, instructions may access any region in the local virtual address space, when a LTB entry matches, and may access regions in the global virtual address space when no LTB entry matches. This mechanism permits privileged code to make judicious use of local virtual address ranges, which simplifies the manner in which privileged code may manipulate the contents of a local virtual address range on behalf of a less-privileged client. Note, however, that under this model, an LTB miss does not cause an exception directly, so the use of more local virtual address ranges than LTB entries requires more care: the local virtual address ranges should be selected so as not to overlap with the global virtual address ranges, and GTB misses to LVA regions must be detected and cause the handler to load an LTB entry.

Each thread has an independent LTB, so that threads may independently define local translation. The size of the LTB for each thread is implementation dependent and defined as the **LE** parameter in the architecture description register. **LE** is the log of the number of entries in the local TB per thread; an implementation may define **LE** to be a minimum of 0, meaning one LTB entry per thread, or a maximum of 3, meaning eight LTB entries per thread. For the initial Zeus implementation, each thread has two entries and **LE=1**.


A minimum implementation of an LTB context is a single set of lm/la/lx/lp registers per thread. However, the need for the LTB to translate both code addresses and data addresses

imposes some limits on the use of the LTB in such systems. We need to be able to guarantee forward progress. With a single LTB set per thread, either the code or the data must use global addresses, or both must use the same local address range, as must the LTB and GTB exception handler. To avoid this restriction, the implementation must be raised to two sets per thread, at least one for code and one for data, to guarantee forward progress for arbitrary use of local addresses in the user code (but still be limited to using global addresses for exception handlers).

A single-set LTB context may be further simplified by reserving the implementation of the lm and la registers, setting them to a read-only zero value: Note that in such a configuration, only a single LA region can be implemented.




If the largest possible space is reserved for an address space identifier, the virtual address is partitioned as shown below. Any of the bits marked as "local" below may be used as "offset" as desired.



To improve performance, an implementation may perform the LTB translation on the value of the base register (rc) or unincremented program counter, provided that a check is performed which prohibits changing the unmasked upper 16 bits by the add or increment. If this optimization is provided and the check fails, an AccessDisallowedByVirtualAddress should be signaled. If this optimization is provided, the architecture description parameter LB=1. Otherwise LTB translation is performed on the local address, la, no checking is required, and LB=0.

The LTB protect field controls the minimum privilege level required for each memory action of read (r), write (w), execute (x), and gateway (g), as well as memory and cache attributes of write allocate (wa), detail access (da), strong ordering (so), cache disable (cd), and write through (wt). These fields are combined with corresponding bits in the GTB protect field to control these attributes for the mapped memory region.



Field Description

The meaning of the fields are given by the following table:

name	size	meaning
g	2	minimum privilege required for gateway access
x	2	minimum privilege required for execute access
w	2	minimum privilege required for write access
r	2	minimum privilege required for read access
0	1	reserved
da	1	detail access
so	1	strong ordering
cc	3	cache control

Definition

```
def ga,LocalProtect ← LocalTranslation(th,ba,la,pl) as
    if LB & {ba63..48 ⊕ la63..48} then
        raise AccessNotAllowedByVirtualAddress
    endif
    me ← NONE
    for i ← 0 to {1 11 0LE}-1
        if {la63..48 & -LocalTB[th][i]63..48} = LocalTB[th][i]47..32 then
            me ← i
        endif
    endfor
    if me = NONE then
        if -ControlRegisterpl+8 then
            raise LocalTBMiss
        endif
        ga ← la
        LocalProtect ← 0
    else
        ga ← {va63..48 ^ LocalTB[th][me]31..16} 11 va47..0
        LocalProtect ← LocalTB[th][me]15..0
    endif
enddef
```

Global Translation Buffer

Global virtual addresses which fail to be accessed in either the LZC, the MTB, the BTB, or PTB are translated to physical references in a table, here named the “Global Translation Buffer,” (GTB).

Each processor may have one or more GTB’s, with each GTB shared by one or more threads. The parameter **GT**, the base-two log of the number of threads which share a GTB, and the parameter **T**, the number of threads, allow computation of the number of GTBs ($T/2^{GT}$), and the number of threads which share each GTB (2^{GT}).

If there are two GTBs and four threads ($GT=1$, $T=4$), GTB 0 services references from threads 0 and 1, and GTB 1 services references from threads 2 and 3.


In the first implementation, there is one GTB, shared by all four threads ($GT=2$, $T=4$). The GTB has 128 entries ($G=7$).

Per clock cycle, each GTB can translate one global virtual address to a physical address, yielding protection information as a side effect.

A GTB miss causes a software trap. This trap is designed to permit a fast handler for GlobalTBMiss to be written in software, by permitting a second GTB miss to occur as an exception, rather than a machine check.

Physical address

There may be as many GTB as threads, and up to 2^{15} entries per GTB. The physical address of a GTB entry for thread th , entry en , byte b is:




Note that in the diagram above, the low-order GT bits of the th value are ignored, reflecting that 2^{GT} threads share a single GTB. A single GTB shared between threads appears multiple times in the address space. GTB entries are packed together so that entries in a GTB are consecutive:

Definition

```
def data,flags ← AccessPhysicalGTB(pa,op,wdata) as
    th ← pa23..19+GT || 0GT
    en ← pa18..4
    if (en < (1 || 0G)) and (th < T) and (pa18+GT..19 = 0) then
        case op of
            R:
                data ← GTBArray[th5..GT][en]
            W:
                GTBArray[th5..GT][en] ← wdata
        endcase
    else
        data ← 0
    endif
enddef
```

Entry Format

Each GTB entry is 128 bits. The format of a GTB entry is:

Field Description

gs = ga + size/2: $256 \leq \text{size} \leq 2^{64}$, ga, global address, is aligned (a multiple of) size.

px = pa ^ ga. pa, ga, and px are all aligned (a multiple of) size

The meaning of the fields are given by the following table:

name	size	meaning
gs	57	global address with size
px	56	physical xor
g	2	minimum privilege required for gateway access
x	2	minimum privilege required for execute access
w	2	minimum privilege required for write access
r	2	minimum privilege required for read access
0	1	reserved
da	1	detail access
so	1	strong ordering
cc	3	cache control

If the entire contents of the GTB entry is zero (0), the entry will not match any global address at all. If a zero value is written, a zero value is read for the GTB entry. Software must not write a zero value for the gs field unless the entire entry is a zero value.

It is an error to write GTB entries that multiply match any global address; all GTB entries must have unique, non-overlapping coverage of the global address space. Hardware may produce a machine check if such overlapping coverage is detected, or may produce any physical address and protection information and continue execution.

Limiting the GTB entry size to 128 bits allows up to replace entries atomically (with a single store operation), which is less complex than the previous design, in which the mask portion was first reduced, then other entries changed, then the mask is expanded. However, it is limiting the amount of attribute information

or physical address range we can specify. Consequently, we are encoding the size as a single additional bit to the global address in order to allow for attribute information.

Definition

```
def pa,GlobalProtect ← GlobalAddressTranslation(th,ga,pl,lda) as
    me ← NONE
    for i ← 0 to (1 11 0G) -1
        if GlobalTB[th5..GT][i] ≠ 0 then
            size ← (GlobalTB[th5..GT][i]63..7 and (064-GlobalTB[th5..GT][i]63..7)) 11 08
            if ((ga63..81108) ^ (GlobalTB[th5..GT][i]63..81108)) and (064.size) = 0 then
                me ← GlobalTB[th5..GT][i]
            endif
        endif
    endfor
    if me = NONE then
        if lda then
            PerformAccessDetail|AccessDetailRequiredByLocalTB)
        endif
        raise GlobalTBMiss
    else
        pa ← (ga63..8 ^ GlobalTB[th5..GT][me]127..72) 11 ga7..0
        GlobalProtect ← GlobalTB[th5..GT][me]71..64 11 01 11 GlobalTB[th5..GT][me]6..0
    endif
enddef
```

GTB Registers

Because the processor contains multiple threads of execution, even when taking virtual memory exceptions, it is possible for two threads to nearly simultaneously invoke software GTB miss exception handlers for the same memory region. In order to avoid producing improper GTB state in such cases, the GTB includes access facilities for indivisibly checking and then updating the contents of the GTB as a result of a memory write to specific addresses.

A 128-bit write to the address GTBUpdateFill (fill=1), as a side effect, causes first a check of the global address specified in the data against the GTB. If the global address check results in a match, the data is directed to write on the matching entry. If there is no match, the address specified by GTBLast is used, and GTBLast is incremented. If incrementing GTBLast results in a zero value, GTBLast is reset to GTBFirst, and GTBBump is set. Note that if the size of the updated value is not equal to the size of the matching entry, the global address check may not adequately ensure that no other entries also cover the address range of the updated value. The operation is unpredictable if multiple entries match the global address.

The GTBUpdateFill register is a 128-bit memory-mapped location, to which a write operation performs the operation defined above. A read operation returns a zero value. The format of the GTBUpdateFill register is identical to that of a GTB entry.


An alternative write address, GTBUpdate, (fill=0) updates a matching entry, but makes no change to the GTB if no entry matches. This operation can be used to indivisibly update a GTB entry as to protection or physical address information.

Definition

```
def GTBUpdateWrite(th,fill,data) as
    me ← NONE
    for i ← 0 to (1 11 0G) -1
        size ← (GlobalTB[th5..GT][i]63..7 and (064-GlobalTB[th5..GT][i]63..7)) 11 08
        if ((data63..81108) ^ (GlobalTB[th5..GT][i]63..81108) and (064-size) = 0 then
            me ← i
        endif
    endfor
    if me = NONE then
        if fill then
            GlobalTB[th5..GT][GTBLast[th5..GT]] ← data
            GTBLast[th5..GT] ← (GTBLast[th5..GT] + 1)G-1..0
            if GTBLast[th5..GT] = 0 then
                GTBLast[th5..GT] ← GTBFirst[th5..GT]
                GTBBump[th5..GT] ← 1
            endif
        endif
    else
        GlobalTB[th5..GT][me] ← data
    endif
enddef .
```

Physical address


There may be as many GTB as threads, and up to 2^{11} registers per GTB (5 registers are implemented). The physical address of a GTB control register for thread th, register rn, byte b is:



Note that in the diagram above, the low-order GT bits of the th value are ignored, reflecting that 2^G threads share single GTB registers. A single set of GTB registers shared between threads appears multiple times in the address space, and manipulates the GTB of the threads with which the registers are associated.

The GTBUpdate register is a 128-bit memory-mapped location, to which a write operation performs the operation defined above. A read operation returns a zero value. The format of the GTBUpdateFill register is identical to that of a GTB entry.

The registers GTBLast, GTBFirst, and GTBBump are memory mapped. The GTBLast and GTBFirst registers are G bits wide, and the GTBBump register is one bit.:

Definition

```

def data,flags ← AccessPhysicalGTBRegisters(pa,op,wdata) as
    th ← pa23..19+GT || 0GT
    rn ← pa18..8
    if (rn < 5) and (th < T) and (pa18+GT..19 = 0) and (pa7..4 = 0) then
        case rn || op of
            0 || R, 1 || R:
                data ← 0
            0 || W, 1 || W:
                GTBUpdateWrite[th,rn0,wdata]
            2 || R:
                data ← 064-G || GTBLast[th5..GT]
            2 || W:
                GTBLast[th5..GT] ← wdataG-1..0
            3 || R:
                data ← 064-G || GTBFist[th5..GT]
            3 || W:
                GTBFist[th5..GT] ← wdataG-1..0
            3 || R:
                data ← 063 || GTBBump[th5..GT]
            3 || W:
                GTBBump[th5..GT] ← wdata0
        endcase
    else
        data ← 0
    endif
enddef

```

Address Generation


The address units of each of the four threads provide up to two global virtual addresses of load, store, or memory instructions, for a total of eight addresses. LTB units associated with each thread translate the local addresses into global addresses. The LZC operates on global addresses. MTB, BTB, and PTB units associated with each thread translate the global addresses into physical addresses and cache addresses. (A PTB unit associated with each thread produces physical addresses and cache addresses for program counter references. – this is optional, as by limiting address generation to two per thread, the MTB can be used for

program references.) Cache addresses are presented to the LOC as required, and physical addresses are checked against cache tags as required.

Memory Banks


The LZC has two banks, each servicing up to four requests. The LOC has eight banks, each servicing at most one request.

Assuming random request addresses, the following graph shows the expected rate at which requests are serviced by multi-bank/multi-port memories that have 8 total ports and divided into 1, 2, 4, or 8 interleaved banks. The LZC is 2 banks, each with 4 ports, and the LOC is 8 banks, each 1 port.



Note a small difference between applying 12 references versus 8 references for the LOC (6.5 vs 5.2), and for the LZC (7.8 vs. 6.9). This suggests that simplifying the system to produce two address per thread (program+load/store or two load/store) will not overly hurt performance. A closer simulation, taking into account the sequential nature of the program and load/store traffic may well yield better numbers, as threads will tend to line up in non-interfering patterns, and program microcaching reduces program fetching.

The following graph shows the rates for both 8 total ports and 16 total ports.



Note significant differences between 8-port systems and 16-port systems, even when used with a maximum of 8 applied references. In particular, a 16-bank 1-port system is better than a 4-bank 2-port system with more than 6 applied references. Current layout estimates would require about a 14% area increase (assuming no savings from smaller/simpler sense amps) to switch to a 16-port LOC, with a 22% increase in 8-reference throughput.

Program Microcache

A program microcache (PMC) which holds only program code for each thread may optionally exist, and does exist for the initial implementation. The program microcache is flushed by `reset`, or by executing a `B.BARRIER` instruction. The program microcache is always clean, and is not snooped by writes or otherwise kept coherent, except by flushing as indicated above. The microcache is not altered by writing to the LTB or GTB, and software must execute a `B.BARRIER` instruction before expecting the new contents of the LTB or GTB to affect determination of PMC hit or miss status on program fetches.

In the initial implementation, the program microcache holds simple loop code. The microcache holds two separately addressed cache lines. Branches or execution beyond this region cause the microcache to be flushed and refilled at the new address, provided that the addresses are executable by the current thread. The program microcache uses the B.HINT and B.HINT.I to accelerate fetching of program code when possible. The program microcache generally functions as a prefetch buffer, except that short forward or backward branches within the region covered maintain the contents of the microcache.

Program fetches into the microcache are requested on any cycle in which less than two load/store addresses are generated by the address unit, unless the microcache is already full. System arbitration logic should give program fetches lower priority than load/store references when first presented, then equal priority if the fetch fails arbitration a certain number of times. The delay until program fetches have equal priority should be based on the expected time the program fetch data will be executed; it may be as small as a single cycle, or greater for fetches which are far ahead of the execution point.

Wide Microcache

A wide microcache (WMC) which holds only data fetched for wide (W) instructions may optionally exist, and does exist for the initial implementation, for each unit which implements one or more wide (W) instructions.

The wide (W) instructions each operate on a block of data fetched from memory and the contents of one or more registers, producing a result in a register. Generally, the amount of data in the block exceeds the maximum amount of data that the memory system can supply in a single cycle, so caching the memory data is of particular importance. All the wide (W) instructions require that the memory data be located at an aligned address, an address that is a multiple of the size of the memory data, which is always a power of two.

The wide (W) instructions are performed by functional units which normally perform execute or "back-end" instructions, though the loading of the memory data requires use of the access or "front-end" functional units. To minimize the use of the "front-end" functional units, special rules are used to maintain the coherence of a wide microcache (WMC).

Execution of a wide (W) instruction has a residual effect of loading the specified memory data into a wide microcache (WMC). Under certain conditions, a future wide (W) instruction may be able to reuse the WMC contents.

First of all, any store or cache coherency action on the physical addresses referenced by the WMC will invalidate the contents. The minimum translation unit of the virtual memory system, 256 bytes, defines the number of physical address blocks which must be checked by any store. A WMC for the W.TABLE instruction may be as large as 4096 bytes, and so requires as many as 16 such physical address blocks to be checked for each WMC entry. A WMC for the W.SWITCH or W.MUL.* instructions need check only one address block for each WMC entry, as the maximum size is 128 bytes.

By making these checks on the physical addresses, we do not need to be concerned about changes to the virtual memory mapping from virtual to physical addresses, and the virtual memory state can be freely changed without invalidating any WMC.

Absent any of the above changes, the WMC is only valid if it contains the contents relevant to the current wide (W) instruction. To check this with minimal use of the front-end units, each WMC entry contains a first tag with the thread and address register for which it was last used. If the current wide (W) instruction uses the same thread and address register, it may proceed safely. Any intervening writes to that address register by that thread invalidates the WMC thread and address register tag.

If the above test fails, the front-end is used to fetch the address register and check its contents against a second WMC tag, with the physical addresses for which it was last used. If the tag matches, it may proceed safely. As detailed above, any intervening stores or cache coherency action by any thread to the physical addresses invalidates the WMC entry.

If both the above tests fail for all relevant WMC entries, there is no alternative but to load the data from the virtual memory system into the WMC. The front-end units are responsible for generating the necessary addresses to the virtual memory system to fetch the entire data block into a WMC.

For the first implementation, it is anticipated that there be eight WMC entries for each of the two X units (for W.SWITCH instructions), eight WMC entries for each of the two E units (for W.MUL instructions), and four WMC entries for the single T unit. The total number of WMC address tags requires is $8*2*1+8*2*1+4*1*16 = 96$ entries.

The number of WMC address tags can be substantially reduced to $32+4=36$ entries by making an implementation restriction requiring that a single translation block be used to translate the data address of W.TABLE instructions. With this restriction, each W.TABLE WMC entry uses a contiguous and aligned physical data memory block, for which a single address tag can contain the relevant information. The size of such a block is a maximum of 4096 bytes. The restriction can be checked by examining the size field of the referenced GTB entry.

Level Zero Cache

The innermost cache level, here named the "Level Zero Cache," (LZC) is fully associative and indexed by global address. Entries in the LZC contain global addresses and previously fetched data from the memory system. The LZC is an implementation feature, not visible to the Zeus architecture.

Entries in the LZC are also used to hold the global addresses of store instructions that have been issued, but not yet completed in the memory system. The LZC entry may also contain the data associated with the global address, as maintained either before or after updating with the store data. When it contains the post-store data, results of stores may be forwarded directly to the requested reference.

With an LZC hit, data is returned from the LZC data, and protection from the LZC tag. No LOC access is required to complete the reference.

All loads and program fetches are checked against the LZC for conflicts with entries being used as store buffer. On a LZC hit on such entries, if the post-store data is present, data may be returned by the LZC to satisfy the load or program fetch. If the post-store data is not present, the load or program fetch must stall until the data is available.

With an LZC miss, a victim entry is selected, and if dirty, the victim entry is written to the LOC. The LOC cache is accessed, and a valid LZC entry is constructed from data from the LOC and tags from the LOC protection information.

All stores are checked against the LZC for conflicts, and further cause a new entry in the LZC, or "take over" a previously clean LZC entry for this purpose. Unaligned stores may require two entries in the LZC. At time of allocation, the address is filled in.

Two operations then occur in parallel - 1) for write-back cached references, the remaining bytes of the hexlet are loaded from the LOC (or LZC), and 2) the addressed bytes are filled in with data from data path. If an exception causes the store to be purged before retirement, the LZC entry is marked invalid, and not written back. When the store is retired, the LZC entry can be written back to LOC or external interface.

Structure

The eight memory addresses are partitioned into up to four odd addresses, and four even addresses.

The LZC contains 16 fully associative entries that may each contain a single hexlet of data at even hexlet addresses (LZCE), and another 16 entries for odd hexlet addresses (LZCO). The maximum capacity of the LZC is $16*32=512$ bytes.

The tags for these entries are indexed by global virtual address (63..5), and contain access control information, detailed below.

The address of entries accessed associatively is also encoded into binary and provided as output from the tags for use in updating the LZC, through its write ports.

8 bit rwxg

16 bit valid

16 bit dirty

4 bit L0\$ address

16 bit protection

```

def data,protect,valid,dirty,match ← LevelZeroCacheRead(ga) as
    eo ← ga4
    match ← NONE
    for i ← 0 to LevelZeroCacheEntries/2-1
        if (ga63..5 = LevelZeroTag[eo][i] then
            match ← i
        endif
    endfor
    if match = NONE then
        raise LevelZeroCacheMiss
    else
        data ← LevelZeroData[eo][match]127..0
        valid ← LevelZeroData[eo][match]143..128
        dirty ← LevelZeroData[eo][match]159..144
        protect ← LevelZeroData[eo][match]167..160
    endif
enddef

```

Level One Cache


The next cache level, here named the “Level One Cache,” (LOC) is four-set-associative and indexed by the physical address. The eight memory addresses are partitioned into up to eight addresses for each of eight independent memory banks. The LOC has a cache block size of 256 bytes, with triclet (32-byte) sub-blocks.

The LOC may be partitioned into two sections, one part used as a cache, and the remainder used as “niche memory.” Niche memory is at least as fast as cache memory, but unlike cache, never misses to main memory. Niche memory may be placed at any virtual address, and has physical addresses fixed in the memory map. The **nl** field in the **control register** configures the partitioning of LOC into cache memory and niche memory.

The LOC data memory is $(256+8)\times 4 \times (128+2)$ bits, depth to hold 256 entries in each of four sets, each entry consisting of one hexlet of data (128 bits), one bit of parity, and one spare bit. The additional 8 entries in each of four sets hold the LOC tags, with 128 bits per entry for 1/8 of the total cache, using 512 bytes per data memory and 4K bytes total.

There are 128 cache blocks per set, or 512 cache blocks total. The maximum capacity of the LOC is 128k bytes. Used as a cache, the LOC is partitioned into 4 sets, each 32k bytes. Physically, the LOC is partitioned into 8 interleaved physical blocks, each holding 16k bytes.


The physical address **pa63..0** is partitioned as below into a 52 to 54 bit **tag** (three to five bits are duplicated from the following field to accommodate use of portion of the cache as niche), 8-bit address to the memory bank (7 bits are physical address (**pa**), 1 bit is virtual address (**v**)), 3 bit memory bank select (**bn**), and 4-bit byte address (**bt**). All access to the LOC are in units of 128 bits (hexlets), so the 4-bit byte address (**bt**) does not apply here. The shaded field (**pa,v**) is translated via **nl** to a cache identifier (**ci**) and set identifier (**si**) and presented to the LOC as the LOC address to LOC bank **bn**.



The LOC tag consists of 64 bits of information, including a 52 to 54-bit tag and other cache state information. Only one MTB entry at a time may contain a LOC tag.

With 256 byte cache lines, there are 512 cache blocks. At 64 bits per tag, the cache tags require 4k bytes of storage. This storage is adjacent to the LOC data memory itself, using physical addresses = 1024..1055. Alternatively (see detailed description below), physical addresses = 0..31 may be used.

The format of a LOC tag entry is shown below.




The meaning of the fields are given by the following table:

name	size	meaning
tag	52	physical address tag
da	1	detail access (or physical address bit 11)
vs	1	victim select (or physical address bit 10)
mesi	2	coherency: modified (3), exclusive (2), shared (1), invalid (0)
tv	8	triclet valid (1) or invalid (0)

To access the LOC, a global address is supplied to the Micro-Tag Buffer (MTB), which associatively looks up the global address into a table holding a subset of the LOC tags. In particular, each MTB table entry contains the cache index derived from physical address bits 14..8, ci, (7 bits) and set identifier, si, (2 bits) required to access the LOC data. Each MTB table entry also contains the protection information of the LOC tag.

With an MTB hit, protection information is supplied from the MTB. The MTB supplies the resulting cache index (ci, from the MTB), set identifier, si, (2 bits) and virtual address (bit 7, v, from the LA), which are applied to the LOC data bank selected from bits 6..4 of the LA. The diagram below shows the address presented to LOC data bank bn.



With an MTB miss, the GTB (described below) is referenced to obtain a physical address and protection information.

To select the cache line, a 7-bit niche limit register **nl** is compared against the value of **pa14..8** from the GTB. If $\text{pa14..8} < \text{nl}$, a 7-bit address modifier register **am** is inclusive-or'ed against **pa14..8**, producing a cache index, **ci**. Otherwise, **pa14..8** is used as **ci**. Cache lines 0..**nl**-1, and cache tags 0..**nl**-1, are available for use as niche memory. Cache lines **nl**..127 and cache tags **nl**..127 are used as LOC.


$$\text{ci} \leftarrow (\text{pa14..8} < \text{nl}) ? (\text{pa14..8} \text{ I } \text{am}) : \text{pa14..8}$$

The address modifier **am** is $(1^{7-\log(128-\text{nl})} \mid\mid 0^{\log(128-\text{nl})})$. The **bt** field specifies the least-significant bit used for tag, and is $(\text{nl} < 112) ? 12 : 8 + \log(128-\text{nl})$:

nl	am	bt
0	0	12
1..64	64	12
65..96	96	12
97..112	112	12
113..120	120	11
121..124	124	10
125..126	126	9
127	127	8

Values for **nl** in the range 113..127 require more than 52 physical address tag bits in the LOC tag and a requisite reduction in LOC features. Note that the presence of bits 14..10 of the physical address in the LOC tag is a result of the possibility that, with **am**=64..127, the cache index value **ci** cannot be relied upon to supply bit 14..8. Bits 9..8 can be safely inferred from the cache index value **ci**, so long as **nl** is in the range 0..124. When **nl** is in the range 113..127, the **da** bit is used for bit 11 of the physical address, so the Tag detail access bit is suppressed. When **nl** is in the range 121..127, the **vs** bit is used for bit 10 of the physical address, so victim selection is performed without state bits in the LOC tag. When **nl** is in the range 125..127, the set associativity is decreased, so that **si1** is used for bit 9 of the physical address and when **nl** is 127, **si0** is used for bit 8 of the physical address.

Four tags are fetched from the LOC tags and compared against the PA to determine which of the four sets contain the data. The four tags are contained in two consecutive banks; they may be simultaneously or independently fetched. The diagram below shows the address presented to LOC data bank (**ci1..0** || **si1**).



Note that the **CT** architecture description variable is present in the above address. **CT** describes whether dedicated locations exist in the LOC for tags at the next power-of-two boundary above the LOC data. The niche-mapping mechanism can provide the storage for the LOC tags, so the existence of these dedicated tags is optional: If **CT**=0, addresses at the beginning of the LOC (0..31 for this implementation) are used for LOC tags, and the **nl** value should be adjusted accordingly by software.

The LOC address ($ci \mid\mid si$) uniquely identifies the cache location, and this LOC address is associatively checked against all MTB entries on changes to the LOC tags, such as by cache block replacement, bus snooping, or software modification. Any matching MTB entries are flushed, even if the MTB entry specifies a different global address - this permits address aliasing (the use of a physical address with more than one global address).

With an LOC miss, a victim set is selected (LOC victim selection is described below), whose contents, if any sub-block is modified, is written to the external memory. A new LOC entry is constructed with address and protection information from the GTB, and data fetched from external memory.

The diagram below shows the contents of LOC data memory banks 0..7 for addresses 0..2047:

address	bank 7	...	bank 1	bank 0
0	line 0, hexlet 7, set 0		line 0, hexlet 1, set 0	line 0, hexlet 0, set 0
1	line 0, hexlet 15, set 0		line 0, hexlet 9, set 0	line 0, hexlet 8, set 0
2	line 0, hexlet 7, set 1		line 0, hexlet 1, set 1	line 0, hexlet 0, set 1
3	line 0, hexlet 15, set 1		line 0, hexlet 9, set 1	line 0, hexlet 8, set 1
4	line 0, hexlet 7, set 2		line 0, hexlet 1, set 2	line 0, hexlet 0, set 2
5	line 0, hexlet 15, set 2		line 0, hexlet 9, set 2	line 0, hexlet 8, set 2
6	line 0, hexlet 7, set 3		line 0, hexlet 1, set 3	line 0, hexlet 0, set 3
7	line 0, hexlet 15, set 3		line 0, hexlet 9, set 3	line 0, hexlet 8, set 3
8	line 1, hexlet 7, set 0		line 1, hexlet 1, set 0	line 1, hexlet 0, set 0
9	line 1, hexlet 15, set 0		line 1, hexlet 9, set 0	line 1, hexlet 8, set 0
10	line 1, hexlet 7, set 1		line 1, hexlet 1, set 1	line 1, hexlet 0, set 1
11	line 1, hexlet 15, set 1		line 1, hexlet 9, set 1	line 1, hexlet 8, set 1
12	line 1, hexlet 7, set 2		line 1, hexlet 1, set 2	line 1, hexlet 0, set 2
13	line 1, hexlet 15, set 2		line 1, hexlet 9, set 2	line 1, hexlet 8, set 2
14	line 1, hexlet 7, set 3		line 1, hexlet 1, set 3	line 1, hexlet 0, set 3
15	line 1, hexlet 15, set 3		line 1, hexlet 9, set 3	line 1, hexlet 8, set 3
...
1016	line 127, hexlet 7, set 0		line 127, hexlet 1, set 0	line 127, hexlet 0, set 0
1017	line 127, hexlet 15, set 0		line 127, hexlet 9, set 0	line 127, hexlet 8, set 0
1018	line 127, hexlet 7, set 1		line 127, hexlet 1, set 1	line 127, hexlet 0, set 1
1019	line 127, hexlet 15, set 1		line 127, hexlet 9, set 1	line 127, hexlet 8, set 1
1020	line 127, hexlet 7, set 2		line 127, hexlet 1, set 2	line 127, hexlet 0, set 2
1021	line 127, hexlet 15, set 2		line 127, hexlet 9, set 2	line 127, hexlet 8, set 2
1022	line 127, hexlet 7, set 3		line 127, hexlet 1, set 3	line 127, hexlet 0, set 3
1023	line 127, hexlet 15, set 3		line 127, hexlet 9, set 3	line 127, hexlet 8, set 3
1024	tag line 3, sets 3 and 2		tag line 0, sets 3 and 2	tag line 0, sets 1 and 0
1025	tag line 7, sets 3 and 2		tag line 4, sets 3 and 2	tag line 4, sets 1 and 0
...
1055	tag line 127, sets 3 and 2		tag line 124, sets 3 and 2	tag line 124, sets 1 and 0
1056	reserved		reserved	reserved
...
2047	reserved		reserved	reserved

The following table summarizes the state transitions required by the LOC cache:

cc	op	mesi	v	bus op	c	x	mesi	v	w	m	notes
NC	R	x	x	uncached read							
NC	W	x	x	uncached write							
CD	R	I	x	uncached read							
CD	R	x	0	uncached read							
CD	R	MES	1	(hit)							
CD	W	I	x	uncached write							
CD	W	x	0	uncached write							
CD	W	MES	1	uncached write						1	
WT/WA	R	I	x	triclet read	0	x					
WT/WA	R	I	x	triclet read	1	0	S	1			
WT/WA	R	I	x	triclet read	1	1	E	1			
WT/WA	R	MES	0	triclet read	0	x					inconsistent KEN#
WT/WA	R	S	0	triclet read	1	0		1			
WT/WA	R	S	0	triclet read	1	1		1			E->S: extra sharing
WT/WA	R	E	0	triclet read	1	0		1			
WT/WA	R	E	0	triclet read	1	1	S	1			shared block
WT/WA	R	M	0	triclet read	1	0	S	1			other subblocks M->I
WT/WA	R	M	0	triclet read	1	1		1			E->M: extra dirty
WT/WA	R	MES	1	(hit)							
WT	W	I	x	uncached write							
WT	W	x	0	uncached write							
WT	W	MES	1	uncached write					1		
WA	W	I	x	triclet read	0	x		1			throwaway read
WA	W	I	x	triclet read	1	0	S	1	1	1	
WA	W	I	x	triclet read	1	1	M	1	1	1	
WA	W	MES	0	triclet read	0	x		1	1		inconsistent KEN#
WA	W	S	0	triclet read	1	0	S	1	1	1	
WA	W	S	0	triclet read	1	1	M	1	1	1	
WA	W	S	1	write	0	S		1	1		
WA	W	S	1	write	1	S		1	1		E->S: extra sharing
WA	W	E	0	triclet read	1	0	S	1	1	1	
WA	W	E	0	triclet read	1	1	E	1	1	1	
WA	W	E	1	(hit)		x	M	1			E->M: extra dirty
WA	W	M	0	triclet read	1	0	M	1	1	1	
WA	W	M	0	triclet read	1	1	M	1	1	1	
WA	W	M	1	(hit)		x	M	1			

cc	cache control
op	operation: R=read, W=write
mesi	current mesi state
v	current tv state
bus op	bus operation
c	cachable (triclet) result
x	exclusive result
mesi	new mesi state
v	new tv state
w	cacheable write after read
m	merge store data with cache line data
notes	other notes on transition

Definition

```

def data,tda ← LevelOneCacheAccess(pa,size,lda,gda,cc,op,wd) as

    // cache index
    am ← (17-log(128-nl) || 0log(128-nl))
    ci ← (pa14..8<nl) ? (pa14..8||am) : pa14..8
    bt ← (nl≤112) ? 12 : 8+log(128-nl)

    // fetch tags for all four sets
    tag10 ← ReadPhysical(0xFFFFFFF0000000063..19||CT1105||ci||01||04, 128)
    Tag[0] ← tag1063..0
    Tag[1] ← tag10127..64
    tag32 ← ReadPhysical(0xFFFFFFF0000000063..19||CT1105||ci||111||04, 128)
    Tag[2] ← tag3263..0
    Tag[3] ← tag32127..64
    vsc ← (Tag[3]10 || Tag[2]10) ^ (Tag[1]10 || Tag[0]10)

    // look for matching tag
    si ← MISS
    for i ← 0 to 3
        if (Tag[i]63..10 || i1..0 || 07)63..bt = pa63..bt then
            si ← i
        endif
    endfor

    // detail access checking on MISS
    if (si = MISS) and (lda ≠ gda) then
        if gda then
            PerformAccessDetail(AccessDetailRequiredByGlobalTB)
        else
            PerformAccessDetail(AccessDetailRequiredByLocalTB)
        endif
    endif

    // if no matching tag or invalid MESI or no sub-block, perform cacheable read/write
    bd ← (si = MISS) or (Tag[si]9..8 = I) or ((op=W) and (Tag[si]9..8 = S)) or -Tag[si]pa7..5
    if bd then

```

```

if (op=W) and (cc >= WA) and ((si = MISS) or -Tag[si]pa7..5 or (Tag[si]9..8 ≠ S)) then
    data,cen,xen ← AccessPhysical(pa,size,cc,R,0)
    //if cache disabled or shared, do a write through
    if -cen or -xen then
        data,cen,xen ← AccessPhysical(pa,size,cc,W,wd)
    endif
else
    data,cen,xen ← AccessPhysical(pa,size,cc,op,wd)
endif
al ← cen
else
    al ← 0
endif

// find victim set and eject from cache
if al and (si = MISS or Tag[si]9..8 = I) then
    case bt of
        12..11:
            si ← vsc
        10..8:
            gvsc ← gvsc + 1
            si ← (bt≤9) : pa9 : gvsc1^pa11 || (bt≤8) : pa8 : gvsc0^pa10
    endcase
    if Tag[si]9..8 = M then
        for i ← 0 to 7
            if Tag[si]i then
                vca ← 0xFFFFFFFF0000000063..1911011c11s11i2..01104
                vdata ← ReadPhysical(vca, 256)
                vpa ← {Tag[si]63..10 11 si1..0 11 07}63..bt11pabt-1..811i2..01101104
                WritePhysical(vpa, 256, vdata)
            endif
        endfor
    endif
    if Tag[vsc+1]9..8 = I then
        nvsc ← vsc + 1
    elseif Tag[vsc+2]9..8 = I then
        nvsc ← vsc + 2
    elseif Tag[vsc+3]9..8 = I then
        nvsc ← vsc + 3
    else
        case cc of
            NC, CD, WT, WA, PF:
                nvsc ← vsc + 1
            LS, SS:
                nvsc ← vsc //no change
        endif
    endcase
    endif
    tda ← 0
    sm ← 07-pa7..5 11 11 11 0pa7..5
else
    nvsc ← vsc
    tda ← (bt>11) ? Tag[si]11 : 0
    if al then

```

```

sm ← Tag[si]7..1+pa7..5 || 11 || Tag[si]pa7..5-1..0
endif
endif

// write new data into cache and update victim selection and other tag fields
if al then
    if op=R then
        mesi ← xen ? E : S
    else
        mesi ← xen ? M : I TODO
    endif
    case bt of
        12:
            Tag[si] ← pa63..bt || tda || Tag[si^2]10 ^ nvscsi0 || mesi || sm
            Tag[si^1]10 ← Tag[si^3]10 ^ nvsc1^si0
        11:
            Tag[si] ← pa63..bt || Tag[si^2]10 ^ nvscsi0 || mesi || sm
            Tag[si^1]10 ← Tag[si^3]10 ^ nvsc1^si0
        10:
            Tag[si] ← pa63..bt || mesi || sm
    endcase
    dt ← 1
    nca ← 0xFFFFFFFF0000000063..191011c11s11pa7..51104
    WritePhysical(nca, 256, data)
endif

// retrieve data from cache
if ~bd then
    nca ← 0xFFFFFFFF0000000063..191011c11s11pa7..51104
    data ← ReadPhysical(nca, 128)
endif

// write data into cache
if (op=W) and bd and al then
    nca ← 0xFFFFFFFF0000000063..191011c11s11pa7..51104
    data ← ReadPhysical(nca, 128)
    mdata ← data127..8*(size+pa3..0) || wd8*(size+pa3..0)-1..8*pa3..0 || data8*pa3..0..0
    WritePhysical(nca, 128, mdata)
endif

// prefetch into cache
if al=bd and (cc=PF or cc=LS) then
    af ← 0 // abort fetch if af becomes 1
    for i ← 0 to 7
        if ~Tag[si]i and ~af then
            data,cen,xen ← AccessPhysical(pa63..811i2..01101104,256,cc,R,0)
            if cen then
                nca ← 0xFFFFFFFF0000000063..191011c11s11i2..01104
                WritePhysical(nca, 256, data)
                Tag[si]i ← 1
                dt ← 1
            else

```

```

        af ← 1
    endif
endif
endfor
endif


// cache tag writeback if dirty
if dt then
    nt ← Tag[si11111] || Tag[si11101]
    WritePhysical(0xFFFFFFFF000000063..19||CT1105||ci1|si11104, 128, nt)
endif
enddef

```


Physical address

The LOC data memory banks are accessed implicitly by cached memory accesses to any physical memory location as shown above. The LOC data memory banks are also accessed explicitly by uncached memory accesses to particular physical address ranges. The address mapping of these ranges is designed to facilitate use of a contiguous portion of the LOC cache as niche memory.

The physical address of a LOC hexlet for LOC address **ba**, bank **bn**, byte **b** is:



Within the explicit LOC data range, starting from a physical address **pa17..0**, the diagram below shows the LOC address (**pa17..7**) presented to LOC data bank (**pa6..4**).



The diagram below shows the LOC data memory bank and address referenced by byte address offsets in the explicit LOC data range. Note that this mapping includes the addresses use for LOC tags.

Byte offset	
0	bank 0, address 0
16	bank 1, address 0
32	bank 2, address 0
48	bank 3, address 0
64	bank 4, address 0
80	bank 5, address 0
96	bank 6, address 0
112	bank 7, address 0
128	bank 0, address 1
144	bank 1, address 1
160	bank 2, address 1
176	bank 3, address 1
192	bank 4, address 1
208	bank 5, address 1
224	bank 6, address 1
240	bank 7, address 1
...	
262016	bank 0, address 2047
262032	bank 1, address 2047
262048	bank 2, address 2047
262064	bank 3, address 2047
262080	bank 4, address 2047
262096	bank 5, address 2047
262112	bank 6, address 2047
262128	bank 7, address 2047

Definition

```

def data ← AccessPhysicalLOC(pa,op,wd) as
    bank ← pa6..4
    addr ← pa17..7
    case op of
        R:
            rd ← LOCArray[bank][addr]
            crc ← LOCRedundancy[bank]
            data ← (crc and rd130..2) or (~crc and rd128..0)
            p[0] ← 0
            for i ← 0 to 128 by 1
                p[i+1] ← p[i] ^ data;
            endfor
            if ControlRegister61 and (p[129] ≠ 1) then
                raise CacheError
            endif
        W:
            p[0] ← 0
            for i ← 0 to 127 by 1
                p[i+1] ← p[i] ^ wd;
            endfor
            wd128 ← ~p[128]
            crc ← LOCRedundancy[bank]
            rdata ← (crc126..0 and wd126..0) or (~crc126..0 and wd128..2)

```

```

LOCArray[bank][addr] ← wd128..127 || rdata || wd1..0
endcase
enddef

```

Level One Cache Stress Control

LOC cells may be fabricated with marginal parameters, for which changes in clock timing or power supply voltage may cause these LOC cells to fail or pass. When testing the LOC while the part is in a normal circuit environment, rather than a special test environment with changeable power supply levels, cells with marginal parameters may not reliably fail testing.

To combat this problem, two bits of the control register, **LOC stress**, may be set to stress the circuit environment while testing. Under normal operation, these bits are cleared (00), while during stress testing, one or more of these bits are set (01, 10, 11). Self-testing should be performed in each of the environment settings, and the detected failures combined together to produce a reliable test for cells with marginal parameters.

Level One Cache Redundancy

The LOC contains facilities that can be used to avoid minor defects in the LOC data array.

Each LOC bank has three additional bits of data storage for each 128 bits of memory data (for a total of 131 bits). One of these bits is used to retain odd parity over the 128 bits of memory data, and the other two bits are spare, which can be pressed into service by setting a non-zero value in the LOC redundancy control register for that bank.

Each row of a LOC bank contains 131 bits: 128 bits of memory data, one bit for parity, and two spare bits:

130	129	128	127	0
spare	p		data	
2	1		128	

LOC redundancy control has 129 bits::

128	127	0
pc		control
1		128

Each bit set in the control word causes the corresponding data bit to be selected from a bit address increased by two:

output ← (data and \sim control) or ((spare₀ || p || data_{127..2}) and control)

parity ← (p and \sim pc) or (spare₁ and pc)

The LOC redundancy control register has 129 bits, but is written with a 128-bit value. To set the pc bit in the LOC redundancy control, a value is written to the control with either bit

124 set (1) or bit 126 set (1). To set bit 124 of the LOC redundancy control, a value is written to the control with both bit 124 set (1) and 126 set (1). When the LOC redundancy control register is read, the process is reversed by selecting the pc bit instead of control bit 124 for the value of bit 124 if control bit 126 is zero (0).


This system can remove one defective column at an even bit position and one defective column at an odd bit position within each LOC block. For each defective column location, x, LOC control bit must be set at bits x, x+2, x+4, x+6, ... If the defective column is in the parity location (bit 128), then set bit 124 only. The following table defines the control bits for parity, bit 126 and bit 124: (other control bits are same as values written)

value₁₂₆	value₁₂₄	pc	control₁₂₆	control₁₂₄
0	0	0	0	0
0	1	1	0	0
1	0	1	1	0
1	1	1	1	1

Physical address

The LOC redundancy controls are accessed explicitly by uncached memory accesses to particular physical address ranges.

The physical address of a LOC redundancy control for LOC bank **bn**, byte **b** is:



Definition

```
def data ← AccessPhysicalLOCRedundancy(pa,op,wd) as
    bank ← pa6..4
    case op of
        R:
            rd ← LOCRedundancy[bank]
            data ← rd127..125 || (rd126 ? rd124 : rd128) || rd123..0
        W:
            rd ← (wd126 or wd124) || wd127..125 || (wd126 and wd124) || wd123..0
            LOCRedundancy[bank] ← rd
    endcase
enddef
```

Memory Attributes

Fields in the LTB, GTB and cache tag control various attributes of the memory access in the specified region of memory. These include the control of cache consultation, updating, allocation, prefetching, coherence, ordering, victim selection, detail access, and cache prefetching.

Cache Control

The cache may be used in one of five ways, depending on a three-bit cache control field (cc) in the LTB and GTB. The cache control field may be set to one of seven states: NC, CD, WT, WA, PF, SS, and LS:

State		read		write		read/write	
		consult	allocate	update	allocate	victim	prefetch
No Cache	0	No	No	No	No	No	No
Cache Disable	1	Yes	No	Yes	No	No	No
Write Through	2	Yes	Yes	Yes	No	No	No
reserved	3						
Write Allocate	4	Yes	Yes	Yes	Yes	No	No
PreFetch	5	Yes	Yes	Yes	Yes	No	Yes
SubStream	6	Yes	Yes	Yes	Yes	Yes	No
LineStream	7	Yes	Yes	Yes	Yes	Yes	Yes

The Zeus processor controls cc as an attribute in the LTB and GTB, thus software may set this attribute for certain address ranges and clear it for others. A three-bit field indicates the choice of caching, according to the table above. The maximum of the three-bit cache control field (cc) values of the LTB and GTB indicates the choice of caching, according to the table above.

No Cache

No Cache (NC) is an attribute that can be set on a LTB or GTB translation region to indicate that the cache is to be not to be consulted. No changes to the cache state result from reads or writes with this attribute set, (except for accesses that directly address the cache via memory-mapped region).

Cache Disable

Cache Disable (CD) is an attribute that can be set on a LTB or GTB translation region to indicate that the cache is to be consulted and updated for cache lines which are already present, but no new cache lines or sub-blocks are to be allocated when the cache does not already contain the addressed memory contents.

The "Socket 7" bus also provides a mechanism for supporting chip sets to decide on each access whether data is to be cached, using the CACHE# and KEN# signals. Using these signals, external hardware may cause a region selected as WT, WA or PF to be treated as CD. This mechanism is only active on the first such access to a memory region if caching is enabled, as the cache may satisfy subsequent references without a bus transaction.

Write Through

Write Through (WT) is an attribute that can be set on a LTB or GTB translation region to indicate that the writes to the cache must also immediately update backing memory. Reads to addressed memory that is not present in the cache cause cache lines or sub-blocks to be

allocated. Writes to addressed memory that is not present in the cache does not modify cache state.

The "Socket 7" bus also provides a mechanism for supporting chip sets to decide on each access whether data is to be written through, using the PWT and WB/WT# signals. Using these signals, external hardware may cause a region selected as **WA** or **PF** to be treated as **WT**. This mechanism is only active on the first write to each region of memory; as on subsequent references, if the cache line is in the Exclusive or Modified state and writeback caching is enabled on the first reference, no subsequent bus operation occurs, at least until the cache line is flushed.

Write Allocate

Write allocate (**WA**) is an attribute that can be set of a LTB or GTB translation region to indicate that the processor is to allocate a memory block to the cache when the data is not previously present in the cache and the operation to be performed is a store. Reads to addressed memory that is not present in the cache cause cache lines or sub-blocks to be allocated. For cacheable data, write allocate is generally the preferred policy, as allocating the data to the cache reduces further bus traffic for subsequent references (loads or stores) or the data. Write allocate never occurs for data which is not cached. A write allocate brings in the data immediately into the Modified state.

Other "socket 7" processors have the ability to inhibit write allocate to cached locations under certain conditions, related by the address range. K6, for example, can inhibit write allocate in the range of 15-16Mbyte, or for all addresses above a configurable limit with 4Mbyte granularity. Pentium has the ability to label address ranges over which write allocate can be inhibited.

Prefetch

Prefetch (**PF**) is an attribute that can be set on a LTB or GTB translation region to indicate that increased prefetching is appropriate for references in this region. Each program fetch, load or store to a cache line that does not already contain all the sub-blocks causes a prefetch allocation of the remaining sub-blocks. Cache misses cause allocation of the requested sub-block and prefetch allocation of the remaining sub-blocks. Prefetching does not necessarily fill in the entire cache line, as prefetch memory references are performed at a lower priority to other cache and memory reference traffic. A limited number of prefetches (as low as one in the initial implementation) can be queued; the older prefetch requests are terminated as new ones are created.

In other respects, the **PF** attribute is handled in the manner of the **WA** attribute. Prefetching is considered an implementation-dependent feature, and an implementation may choose to implement region with the **PF** attribute exactly as with the **WA** attribute.

Implementations may perform even more aggressive prefetching in future versions. Data may be prefetched into the cache in regions that are cacheable, as a result of program fetches, loads or stores to nearby addresses. Prefetches may extend beyond the cache line associated with the nearby address. Prefetches shall not occur beyond the reach of the GTB

entry associated with the nearby address. Prefetching is terminated if an attempted cache fill results in a bus response that is not cacheable. Prefetches are implementation-dependent behavior, and such behavior may vary as a result of other memory references or other bus activity.

SubStream

SubStream (SS) is an attribute that can be set on a LTB or GTB translation region to indicate that references in this region are to be selected as the next victim on a cache miss. In particular, cache misses, which normally place the cache line in the last-to-be-victim state, instead place the cache line in the first-to-be-victim state, except relative to cache lines in the I state.

In other respects, the SS attribute is handled in the manner of the WA attribute. SubStream is considered an implementation-dependent feature, and an implementation may choose to implement region with the SS attribute exactly as with the WA attribute.

The SubStream attribute is appropriate for regions which are large data structures in which the processor is likely to reference the memory data just once or a small number of times, but for which the cache permits the data to be fetched using burst transfers. By making it a priority for victimization, these references are less likely to interfere with caching of data for which the cache performs a longer-term storage function.

LineStream

LineStream (LS) is an attribute that can be set on a LTB or GTB translation region to indicate that references in this region are to be selected as the next victim on a cache miss, and to enable prefetching. In particular, cache misses, which normally place the cache line in the last-to-be-victim state, instead place the cache line in the first-to-be-victim state, except relative to cache lines in the I state.

In other respects, the LS attribute is handled in the manner of the PF attribute. LineStream is considered an implementation-dependent feature, and an implementation may choose to implement region with the SS attribute exactly as with the PF or WA attributes.

Like the SubStream attribute, the LineStream attribute is particularly appropriate for regions for which large data structures are used in sequential fashion. By prefetching the entire cache line, memory traffic is performed as large sequential bursts of at least 256 bytes, maximizing the available bus utilization.

Cache Coherence

Cache coherency is maintained by using MESI protocols, for which each cache line (256 bytes) the cache data is kept in one of four states: M, E, S, I:

State		this Cache data	other Cache data	Memory data
Modified	3	Data is held exclusively in this cache.	No data is present in other caches.	The contents of main memory are now invalid.
Exclusive	2	Data is held exclusively in this cache.	No data is present in other caches.	Data is the same as the contents of main memory
Shared	1	Data is held in this cache, and possibly others.	Data is possibly in other caches.	Data is the same as the contents of main memory.
Invalid	0	No data for this location is present in the cache.	Data is possibly in other caches.	Data is possibly present in main memory.

The state is contained in the **mesi** field of the cache tag.

In addition, because the “Socket 7” bus performs block transfers and cache coherency actions on triclet (32 byte) blocks, each cache line also maintains 8 bits of triclet valid (**tv**) state. Each bit of **tv** corresponds to a triclet sub-block of the cache line; bit 0 for bytes 0..31, bit 1 for bytes 32..63, bit 2 for bytes 64..95, etc. If the **tv** bit is zero (0), the coherence state for that triclet is I, no matter what the value of the **mesi** field. If the **tv** bit is one (1), the coherence state is defined by the **mesi** field. If all the **tv** bits are cleared (0), the **mesi** field must also be cleared, indicating an invalid cache line.

Cache coherency activity generally follows the protocols defined by the “Socket 7” bus, as defined by Pentium and K6-2 documentation. However, because the coherence state of a cache line is represented in only 10 bits per 256 bytes (1.25 bits per triclet), a few state transitions are defined differently. The differences are a direct result of attempts to set triclets within a cache line to different MES states that cannot be represented. The data structure allows any triclet to be changed to the I state, so state transitions in this direction match the Pentium processor exactly.

On the Pentium processor, for a cache line in the M state, an external bus Inquiry cycle that does not require invalidation (**INV=0**) places the cache line in the S state. On the Zeus processor, if no other triclet in the cache line is valid, the **mesi** field is changed to S. If other triclets in the cache line are valid, the **mesi** field is left unchanged, and the **tv** bit for this triclet is turned off, effectively changing it to the I state.

On the Pentium processor, for a cache line in the E state, an external bus Inquiry cycle that does not require invalidation (**INV=0**) places the cache line in the S state. On the Zeus processor, the **mesi** field is changed to S. If other triclets in the cache line are valid, the MESI state is effectively changed to the S state for these other triclets.

On the Pentium processor, for a cache line in the S state, an internal store operation causes a write-through cycle and a transition to the E state. On the Zeus processor, the **mesi** field is changed to E. Other triclets in the cache line are invalidated by clearing the tv bits; the MESI state is effectively changed to the I state for these other triclets.

When allocating data into the cache due to a store operation, data is brought immediately into the Modified state, setting the **mesi** field to M. If the previous **mesi** field is S, other triclets which are valid are invalidated by clearing the tv bits. If the previous **mesi** field is E, other triclets are kept valid and therefore changed to the M state.

When allocating data into the cache due to a load operation, data is brought into the Shared state, if another processor reports that the data is present in its cache or the **mesi** field is already set to S, the Exclusive state, if no processor reports that the data is present in its cache and the **mesi** field is currently E or I, or the Modified state if the **mesi** field is already set to M. The determination is performed by driving PWT low and checking whether WB/WT# is sampled high; if so the line is brought into the Exclusive state. (See page 202 (184) of the K6-2 documentation).

Strong Ordering

Strong ordering (**so**) is an attribute which permits certain memory regions to be operated with strong ordering, in which all memory operations are performed exactly in the order specified by the program and others to be operated with weak ordering, in which some memory operations may be performed out of program order.

The Zeus processor controls strong ordering as an attribute in the LTB and GTB, thus software may set this attribute for certain address ranges and clear it for others. A one bit field indicates the choice of access ordering. A one (1) bit indicates strong ordering, while a zero (0) bit indicates weak ordering.

With weak ordering, the memory system may retain store operations in a store buffer indefinitely for later storage into the memory system, or until a synchronization operation to any address performed by the thread that issued the store operation forces the store to occur. Load operations may be performed in any order, subject to requirements that they be performed logically subsequent to prior store operations to the same address, and subsequent to prior synchronization operations to any address. Under weak ordering it is permitted to forward results from a retained store operation to a future load operation to the same address. Operations are considered to be to the same address when any bytes of the operation are in common. Weak ordering is usually appropriate for conventional memory regions, which are side-effect free.

With strong ordering, the memory system must perform load and store operations in the order specified. In particular, strong-ordered load operations are performed in the order specified, and all load operations (whether weak or strong) must be delayed until all previous strong-ordered store operations have been performed, which can have a significant performance impact. Strong ordering is often required for memory-mapped I/O regions, where store operations may have a side-effect on the value returned by loads to other

addresses. Note that Zeus has memory-mapped I/O, such as the TB, for which the use of strong ordering is essential to proper operation of the virtual memory system.

The EWBE# signal in "Socket 7" is of importance in maintaining strong ordering. When a write is performed with the signal inactive, no further writes to E or M state lines may occur until the signal becomes active. Further details are given in Pentium documentation (K6-2 documentation may not apply to this signal.)

Victim Selection

One bit of the cache tag, the vs bit, controls the selection of which set of the four sets at a cache address should next be chosen as a victim for cache line replacement. Victim selection (vs) is an attribute associated with LOC cache blocks. No vs bits are present in the LTB or GTB.

There are two hexlets of tag information for a cache line, and replacement of a set requires writing only one hexlet. To update priority information for victim selection by writing only one hexlet, information in each hexlet is combined by an exclusive-or. It is the nature of the exclusive-or function that altering either of the two hexlets can change the priority information.

Full victim selection ordering for four sets

There are $4 \cdot 3 \cdot 2 \cdot 1 = 24$ possible orderings of the four sets, which can be completely encoded in as few as 5 bits: 2 bits to indicate highest priority, 2 bits for second-highest priority, 1 bit for third-highest priority, and 0 bits for lowest priority. Dividing this up per set and duplicating per hexlet with the exclusive-or scheme above requires three bits per set, which suggests simply keeping track of the three-highest priority sets with 2 bits each, using 6 bits total and three bits per set.

Specifically, vs bits from the four sets are combined to produce a 6-bit value:

$$vsc \leftarrow (vs[3] \mid\mid vs[2]) ^ (vs[1] \mid\mid vs[0])$$

The highest priority for replacement is set vsc_{1..0}, second highest priority is set vsc_{3..2}, third highest priority is set vsc_{5..4}, and lowest priority is vsc_{5..4} ^ vsc_{3..2} ^ vsc_{1..0}. When the highest priority set is replaced, it becomes the new lowest priority and the others are moved up, computing a new vsc by:

$$vsc \leftarrow vsc_{5..4} ^ vsc_{3..2} ^ vsc_{1..0} \mid\mid vsc_{5..2}$$

When replacing set vsc for a LineStream or SubStream replacement, the priority for replacement is unchanged, unless another set contains the invalid MESI state, computing a new vsc by:

$$\begin{aligned} vsc \leftarrow & mesi[vsc_{5..4} ^ vsc_{3..2} ^ vsc_{1..0} = I] ? vsc_{5..4} ^ vsc_{3..2} ^ vsc_{1..0} \mid\mid vsc_{5..2}: \\ & (mesi[vsc_{5..4} = I] ? vsc_{1..0} \mid\mid vsc_{5..2}: \\ & (mesi[vsc_{3..2} = I] ? vsc_{5..4} \mid\mid vsc_{1..0} \mid\mid vsc_{3..2}: \\ & \quad vsc \end{aligned}$$

Cache flushing and invalidations can cause cache lines to be cleared out of sequential order. Flushing or invalidating a cache line moves that set to highest priority. If that set is already highest priority, the vsc is unchanged. If the set was second or third highest or lowest priority, the vsc is changed to move that set to highest priority, moving the others down.

$$vsc \leftarrow ((fs=vsc1..0 \text{ or } fs=vsc3..2) ? vsc5..4 : vsc3..2) || (fs=vsc1..0 ? vsc3..2 : vsc1..0) || fs$$

When updating the hexlet containing vs[1] and vs[0], the new values of vs[1] and vs[0] are:

$$vs[1] \leftarrow vs[3] ^ vsc5..3$$

$$vs[0] \leftarrow vs[2] ^ vsc2..0$$

When updating the hexlet containing vs[3] and vs[2], the new values of vs[3] and vs[2] are:

$$vs[3] \leftarrow vs[1] ^ vsc5..3$$

$$vs[2] \leftarrow vs[0] ^ vsc2..0$$

Software must initialize the vs bits to a legal, consistent state. For example, to set the priority (highest to lowest) to (0, 1, 2, 3), vsc must be set to 0b100100. There are many legal solutions that yield this vsc value, such as vs[3] \leftarrow 0, vs[2] \leftarrow 0, vs[1] \leftarrow 4, vs[0] \leftarrow 4.

Simplified victim selection ordering for four sets

However, the orderings are simplified in the first Zeus implementation, to reduce the number of vs bits to one per set, keeping a two bit vsc state value:

$$vsc \leftarrow (vs[3] || vs[2]) ^ (vs[1] || vs[0])$$

The highest priority for replacement is set vsc, second highest priority is set vsc+1, third highest priority is set vsc+2, and lowest priority is vsc+3. When the highest priority set is replaced, it becomes the new lowest priority and the others are moved up. Priority is given to sets with invalid MESI state, computing a new vsc by:

$$\begin{aligned} vsc \leftarrow & mesi[vsc+1]=I ? vsc + 1 : \\ & (mesi[vsc+2]=I) ? vsc + 2 : \\ & (mesi[vsc+3]=I) ? vsc + 3 : \\ & vsc + 1 \end{aligned}$$

When replacing set vsc for a LineStream or SubStream replacement, the priority for replacement is unchanged, unless another set contains the invalid MESI state, computing a new vsc by:

$$\begin{aligned} vsc \leftarrow & mesi[vsc+1]=I ? vsc + 1 : \\ & (mesi[vsc+2]=I) ? vsc + 2 : \\ & (mesi[vsc+3]=I) ? vsc + 3 : \\ & vsc \end{aligned}$$

Cache flushing and invalidations can cause cache sets to be cleared out of sequential order. If the current highest priority for replacement is a valid set, the flushed or invalidated set is made highest priority for replacement.

$$\text{vsc} \leftarrow (\text{mesi}[\text{vsc}] = 1) ? \text{vsc} : \text{fs}$$

When updating the hexlet containing $\text{vs}[1]$ and $\text{vs}[0]$, the new values of $\text{vs}[1]$ and $\text{vs}[0]$ are:

$$\text{vs}[1] \leftarrow \text{vs}[3] \wedge \text{vsc}_1$$

$$\text{vs}[0] \leftarrow \text{vs}[2] \wedge \text{vsc}_0$$

When updating the hexlet containing $\text{vs}[3]$ and $\text{vs}[2]$, the new values of $\text{vs}[3]$ and $\text{vs}[2]$ are:

$$\text{vs}[3] \leftarrow \text{vs}[1] \wedge \text{vsc}_1$$

$$\text{vs}[2] \leftarrow \text{vs}[0] \wedge \text{vsc}_0$$

Software must initialize the vs bits, but any state is legal. For example, to set the priority (highest to lowest) to (0, 1, 2, 3), vsc must be set to 0b00. There are many legal solutions that yield this vsc value, such as $\text{vs}[3] \leftarrow 0$, $\text{vs}[2] \leftarrow 0$, $\text{vs}[1] \leftarrow 0$, $\text{vs}[0] \leftarrow 0$.

Full victim selection ordering for additional sets

*To extend the full-victim-ordering scheme to eight sets, $3*7=21$ bits are needed, which divided among two tags is 11 bits per tag. This is somewhat generous, as the minimum required is $8*7*6*5*4*3*2*1=40320$ orderings, which can be represented in as few as 16 bits. Extending the full-victim-ordering four-set scheme above to represent the first 4 priorities in binary, but to use 2 bits for each of the next 3 priorities requires $3+3+3+3+2+2+2 = 18$ bits. Representing fewer distinct orderings can further reduce the number of bits used. As an extreme example, using the simplified scheme above with eight sets requires only 3 bits, which divided among two tags is 2 bits per tag.*

Victim selection without LOC tag bits

At extreme values of the **niche limit** register (**nl** in the range 121..124), the bit normally used to hold the vs bit is usurped for use as a physical address bit. Under these conditions, no vsc value is maintained per cache line, instead a single, global vsc value is used to select victims for cache replacement. In this case, the cache consists of four lines, each with four sets. On each replacement a new si value is computed from:

$$\text{gvsc} \leftarrow \text{gvsc} + 1$$

$$\text{si} \leftarrow \text{gvsc} \wedge \text{pa}_{11..10}$$

The algorithm above is designed to utilize all four sets on sequential access to memory.

Victim selection encoding LOC tag bits

At even more extreme values of the **niche limit** register (**nl** in the range 125..127), not only is the bit normally used to hold the **vs** bit usurped for use as a physical address bit, but there is a deficit of one or two physical address bits. In this case, the number of sets can be reduced to encode physical address bits into the victim selection, allowing the choice of set to indicate physical address bits 9 or bits 9..8. On each replacement a new **vsc** value is computed from:

$$\text{gvsc} \leftarrow \text{gvsc} + 1$$

$$\text{si} \leftarrow \text{pa9} \mid\mid (\text{nl}=127) ? \text{pa8} : \text{gvsc}^{\wedge}\text{pa10}$$

The algorithm above is designed to utilize all four sets on sequential access to memory.

Detail Access

Detail access is an attribute which can be set on a cache block or translation region to indicate that software needs to be consulted on each potential access, to determine whether the access should proceed or not. Setting this attribute causes an exception trap to occur, by which software can examine the virtual address, by for example, locating data in a table, and if indicated, causes the processor to continue execution. In continuing, ephemeral state is set upon returning to the re-execution of the instruction that prevents the exception trap from recurring on this particular re-execution only. The ephemeral state is cleared as soon as the instruction is either completed or subject to another exception, so DetailAccess exceptions can recur on a subsequent execution of the same instruction. Alternatively, if the access is not to proceed, execution has been trapped to software at this point, which can abort the thread or take other corrective action.

The detail access attribute permits specification of access parameters over memory region on arbitrary byte boundaries. This is important for emulators, which must prevent store access to code which has been translated, and for simulating machines which have byte granularity on segment boundaries. The detail access attribute can also be applied to debuggers, which have the need to set breakpoints on byte-level data, or which may use the feature to set code breakpoints on instruction boundaries without altering the program code, enabling breakpoints on code contained in ROM.

A one bit field indicates the choice of detail access. A one (1) bit indicates detail access, while a zero (0) bit indicates no detail access. Detail access is an attribute that can be set by the LTB, the GTB, or a cache tag.

The table below indicates the proper status for all potential values of the detail access bits in the LTB, GTB, and Tag:

LTB	GTB	Tag	status
0	0	0	OK - normal
0	0	1	AccessDetailRequiredByTag
0	1	0	AccessDetailRequiredByGTB
0	1	1	OK - GTB inhibited by Tag
1	0	0	AccessDetailRequiredByLTB
1	0	1	OK - LTB inhibited by Tag
1	1	0	OK - LTB inhibited by GTB
1	1	1	AccessDetailRequiredByTag
0	Miss		GTBMiss
1	Miss		AccessDetailRequiredByLTB
0	0	Miss	Cache Miss
0	1	Miss	AccessDetailRequiredByGTB
1	0	Miss	AccessDetailRequiredByLTB
1	1	Miss	Cache Miss

The first eight rows show appropriate activities when all three bits are available. The detail access attributes for the LTB, GTB, and cache tag work together to define whether and which kind of detail access exception trap occurs. Generally, setting a single attribute bit causes an exception, while setting two bits inhibits such exceptions. In this way, a detail access exception can be narrowed down to cause an exception over a specified region of memory: Software generally will set the cache tag detail access bit only for regions in which the LTB or GTB also has a detail access bit set. Because cache activity may flush and refill cache lines implicitly, it is not generally useful to set the cache tag detail access bit alone, but if this occurs, the AccessDetailRequiredByTag exception catches such an attempt.

The next two rows show appropriate activities on a GTB miss. On a GTB miss, the detail access bit in the GTB is not present. If the LTB indicates detail access and the GTB misses, the AccessDetailRequiredByLTB exception should be indicated. If software continues from the AccessDetailRequiredByLTB exception and has not filled in the GTB, the GTBMiss exception happens next. Since the GTBMiss exception is not a continuation exception, a re-execution after the GTBMiss exception can cause a reoccurrence of the AccessDetailRequiredByLTB exception. Alternatively, if software continues from the AccessDetailRequiredByLTB exception and has filled in the GTB, the AccessDetailRequiredByLTB exception is inhibited for that reference, no matter what the status of the GTB and Tag detail bits, but the re-executed instruction is still subject to the AccessDetailRequiredByGTB and AccessDetailRequiredByTag exceptions.

The last four rows show appropriate activities for a cache miss. On a cache miss, the detail access bit in the tag is not present. If the LTB or GTB indicates detail access and the cache misses, the AccessDetailRequiredByLTB or AccessDetailRequiredByGTB exception should be indicated. If software continues from these exceptions and has not filled in the cache, a cache miss happens next. If software continues from the AccessDetailRequiredByLTB or AccessDetailRequiredByGTB exception and has filled in the cache, the previous exception is inhibited for that reference, no matter what the status of the Tag detail bit, but is still subject to the AccessDetailRequiredByTag exception. When the detail bit must be created from a cache miss, the initial value filled in is zero. Software may set the bit, thus turning off AccessDetailRequired exceptions per cache line. If the cache line is flushed and refilled, the

detail access bit in the cache tag is again reset to zero, and another AccessDetailRequired exception occurs.

Settings of the **niche limit** parameter to values that require use of the **da** bit in the LOC tag for retaining the physical address usurp the capability to set the Tag detail access bit. Under such conditions, the Tag detail access bit is effectively always zero (0), so it cannot inhibit AccessDetailRequiredByLTB, inhibit AccessDetailRequiredByGTB, or cause AccessDetailRequiredByTag.

The execution of a Zeus instruction has a reference to one quadlet of instruction, which may be subject to the DetailAccess exceptions, and a reference to data, which may be unaligned or wide. These unaligned or wide references may cross GTB or cache boundaries, and thus involve multiple separate reference that are combined together, each of which may be subject to the DetailAccess exception. There is sufficient information in the DetailAccess exception handler to process unaligned or wide references.

The implementation is free to indicate DetailAccess exceptions for unaligned and wide data references either in combined form, or with each sub-reference separated. For example, in an unaligned reference that crosses a GTB or cache boundary, a DetailAccess exception may be indicated for a portion of the reference. The exception may report the virtual address and size of the complete reference, and upon continuing, may inhibit reoccurrence of the DetailAccess exception for any portion of the reference. Alternatively, it may report the virtual address and size of only a reference portion and inhibit reoccurrence of the DetailAccess exception for only that portion of the reference, subject to another DetailAccess exception occurring for the remaining portion of the reference.


Micro Translation Buffer

The Micro Translation Buffer (MTB) is an implementation-dependent structure which reduces the access traffic to the GTB and the LOC tags. The MTB contains and caches information read from the GTB and LOC tags, and is consulted on each access to the LOC.







To access the LOC, a global address is supplied to the Micro-Translation Buffer (MTB), which associatively looks up the global address into a table holding a subset of the LOC tags. In addition, each table entry contains the physical address bits 14..8 (7 bits) and set identifier (2 bits) required to access the LOC data.

In the first Zeus implementation, there are two MTB blocks - MTB 0 is used for threads 0 and 1, and MTB 1 is used for threads 2 and 3. Per clock cycle, each MTB block can check for 4 simultaneous references to the LOC. Each MTB block has 16 entries.

Each MTB entry consists of a bit less than 128 bits of information, including a 56-bit global address tag, 8 bits of privilege level required for read, write, execute, and gateway access, a detail bit, and 10 bits of cache state indicating for each triclet (32 bytes) sub-block, the MESI state.

MatchOutput

The output of the MTB combines physical address and protection information from the GTB and the referenced cache line.



The meaning of the fields are given by the following table:

name	size	meaning
ga	56	global address
gi	9	GTB index
ci	7	cache index
si	2	set index
vs	12	victim select
da	1	detail access (from cache line)
mesi	2	coherency: modified (3), exclusive (2), shared (1), invalid (0)
tv	8	triclet valid (1) or invalid (0)
g	2	minimum privilege required for gateway access
x	2	minimum privilege required for execute access
w	2	minimum privilege required for write access
r	2	minimum privilege required for read access
0	1	reserved
da	1	detail access (from GTB)
so	1	strong ordering
cc	3	cache control

With an MTB hit, the resulting cache index (14..8 from the MTB, bit 7 from the LA) and set identifier (2 bits from the MTB) are applied to the LOC data bank selected from bits 6..4 of the GVA. The access protection information (pr and rwxg) is supplied from the MTB.

With an MTB (and BTB) miss, a victim entry is selected for replacement. The MTB and BTB are always clean, so the victim entry is discarded without a writeback. The GTB (described below) is referenced to obtain a physical address and protection information. Depending on the access information in the GTB, either the MTB or BTB is filled.

Note that the processing of the physical address $pa_{14..8}$ against the niche limit nl can be performed on the physical address from the GTB, producing the LOC address, ci. The LOC address, after processing against the nl is placed into the MTB directly, reducing the latency of an MTB hit.

Four tags are fetched from the LOC tags and compared against the PA to determine which of the four sets contain the data. If one of the four sets contains the correct physical address, a victim MTB entry is selected for replacement, the MTB is filled and the LOC access proceeds. If none of the four sets is a hit, an LOC miss occurs.

MTB miss GTB cam LOC tag MTB fill

MTB victim

LOC miss

The operation of the MTB is largely not visible to software - hardware mechanisms are responsible for automatically initializing, filling and flushing the MTB. Activity that modifies the GTB or LOC tag state may require that one or more MTB entries are flushed.

A write to the GTBUpdate register that updates a matching entry, a write to the GTBUpdateFill register, or a direct write to the GTB all flush relevant entries from the MTB. MTB flushing is accomplished by searching MTB entries for values that match on the *gi* field with the GTB entry that has been modified. Each such matching MTB entry is flushed.

The MTB is kept synchronous with the LOC tags, particularly with respect to MESI state. On an LOC miss or LOC snoop, any changes in MESI state update (or flush) MTB entries which physically match the address. If the MTB may contain less than the full physical address: it is sufficient to retain the LOC physical address (*ci* || *v* || *si*).

Block Translation Buffer

Zeus has a per thread “Block Translation Buffer” (BTB). The BTB retains GTB information for uncached address blocks. The BTB is used in parallel with the MTB - exactly one of the BTB or MTB may translate a particular reference. When both the BTB and MTB miss, the GTB is consulted, and depending on the result, the block is filled into either the MTB or BTB as appropriate. In the first Zeus implementation, the BTB has 2 entries for each thread.

BTB entries cover any power-of-two granularity, as they retain the size information from the GTB. BTB entries contain no MESI state, as they only contain uncached blocks.

Each BTB entry consists of 128 bits of information, containing the same information in the same format as a GTB entry.

Niche blocks are indicated by GTB information, and correspond to blocks of data that are retained in the LOC and never miss. A special physical address range indicates niche blocks. For this address range, the BTB enables use of the LOC as a niche memory, generating the “set select” address bits from low-order address bits. There is no checking of the LOC tags for consistent use of the LOC as a niche - the *nl* field must be preset by software so that LOC cache replacement never claims the LOC niche space, and only BTB miss and protection bits prevent software from using the cache portion of the LOC as niche.

Other address ranges include other on-chip resources, such as bus interface registers, the control register and status register, as well as off-chip memory, accessed through the bus interface. Each of these regions are accessible as uncached memory.

Program Translation Buffer

Later implementations of Zeus may optionally have a per thread “Program Translation Buffer” (PTB). The PTB retains GTB and LOC cache tag information. The PTB enables generation of LOC instruction fetching in parallel with load/store fetching. The PTB is updated when instruction fetching crosses a cache line boundary (each 64 instructions in straight-line code). The PTB functions similarly to a one-entry MTB, but can use the sequential nature of program code fetching to avoid checking the 56-bit match. The PTB is flushed at the same time as the MTB.

The initial implementation of Zeus has no PTB - the MTB suffices for this function.

Global Virtual Cache

The initial implementation of Zeus contains cache which is both indexed and tagged by a physical address. Other prototype implementations have used a global virtual address to index and/or tag an internal cache. This section will define the required characteristics of a global virtually-indexed cache. TODO

Memory Interface

Dedicated hardware mechanisms are provided to fetch data blocks in the levels zero and one caches, provided that a matching entry can be found in the MTB or GTB (or if the MMU is disabled). Dedicated hardware mechanisms are provided to store back data blocks in the level zero and one caches, regardless of the state of the MTB and GTB. When no entry is to be found in the GTB, an exception handler is invoked either to generate the required information from the virtual address, or to place an entry in the GTB to provide for automatic handling of this and other similarly addressed data blocks.

The initial implementation of Zeus accesses the remainder of the memory system through the "Socket 7" interface. Via this interface, Zeus accesses a secondary cache, DRAM memory, external ROM memory, and an I/O system. The size and presence of the secondary cache and the DRAM memory array, and the contents of the external ROM memory and the I/O system are variables in the processor environment.

Microarchitecture

Each thread has two address generation units, capable of producing two aligned, or one unaligned load or store operation per cycle. Alternatively, these units may produce a single load or store address and a branch target address.

Each thread has a LTB, which translates the two addresses into global virtual addresses.

Each pair of threads has a MTB, which looks up the four references into the LOC. The PTB provides for additional references that are program code fetches.

In parallel with the MTB, these four references are combined with the four references from the other thread pair and partitioned into even and odd hexlet references. Up to four references are selected for each of the even and odd portions of the LZC. One reference for each of the eight banks of the LOC (four are even hexlets; four are odd hexlets) are selected from the eight load/store/branch references and the PTB references.

Some references may be directed to both the LZC and LOC, in which case the LZC hit causes the LOC data to be ignored. An LZC miss which hits in the MTB is filled from the LOC to the LZC. An LZC miss which misses in the MTB causes a GTB access and LOC tag access, then an MTB fill and LOC access, then an LZC fill.

Priority of access: (highest/lowest) cache dump, cache fill, load, program, store.

Snoop

The “Socket 7” bus requires certain bus accesses to be checked against on-chip caches. On a bus read, the address is checked against the on-chip caches, with accesses aborted when requested data is in an internal cache in the M state, and the E state, the internal cache is changed to the S state. On a bus write, data written must update data in on-chip caches. To meet these requirements, physical bus addresses must be checked against the LOC tags.

The S7 bus requires that responses to inquire cycles occur with fixed timing. At least with certain combinations of bus and processor clock rate, inquire cycles will require top priority to meet the inquire response timing requirement.

Synchronization operations must take into account bus activity - generally a synchronization operation can only proceed on cached data which is in Exclusive or Modified – if cached data in Shared state, ownership must be obtained. Data that is not cached must be accessed using locked bus cycles.

Load

Load operations require partitioning into reads that do not cross a hexlet (128 bit) boundary, checking for store conflicts, checking the LZC, checking the LOC, and reading from memory. Execute and Gateway accesses are always aligned and since they are smaller than a hexlet, do not cross a hexlet boundary.

Note: S7 processors perform unaligned operations LSB first, MSB last, up to 64 bits at a time. Unaligned 128 bit loads need 3 64-bit operations, LSB, octlet, MSB. Transfers which are smaller than a hexlet but larger than an octlet are further divided in the S7 bus unit.

Definition

```

def data ← LoadMemoryX{ba,la,size,order)
    assert {order = L} and ({la and (size/8-1)} = 0) and {size = 32}
    hdata ← TranslateAndCacheAccess{ba,la,size,X,0}
    data ← hdata31+8*(la and 15)..8*(la and 15)
enddef

def data ← LoadMemoryG{ba,la,size,order)
    assert {order = L} and ({la and (size/8-1)} = 0) and {size = 64}
    hdata ← TranslateAndCacheAccess{ba,la,size,G,0}
    data ← hdata63+8*(la and 15)..8*(la and 15)
enddef

def data ← LoadMemory{ba,la,size,order)
    if {size > 128} then
        data0 ← LoadMemory{ba, la,size/2, order}
        data1 ← LoadMemory{ba, la+(size/2), size/2, order}
        case order of
            L:
                data ← data1 || data0
            B:

```

```

        data ← data0 || data1
    endcase
else
    bs ← 8*la4..0
    be ← bs + size
    if be > 128 then
        data0 ← LoadMemory(ba, la, 128 - bs, order)
        data1 ← LoadMemory(ba, (la63..5 + 1) || 04, be - 128, order)
        case order of
            L:
                data ← (data1 || data0)
            B:
                data ← (data0 || data1)
        endcase
    else
        hdata ← TranslateAndCacheAccess(ba, la, size, R, 0)
        for i ← 0 to size-8 by 8
            j ← bs + ((order=L) ? i : size-8-i)
            datai+7..i ← hdataj+7..j
        endfor
    endif
endif
enddef

```

Store

Store operations requires partitioning into stores less than 128 bits that do not cross hexlet boundaries, checking for store conflicts, checking the LZC, checking the LOC, and storing into memory.

Definition

```

def StoreMemory(ba, la, size, order, data)
    bs ← 8*la4..0
    be ← bs + size
    if be > 128 then
        case order of
            L:
                data0 ← data127..bs..0
                data1 ← datasize-1..128..bs
            B:
                data0 ← datasize-1..be-128
                data1 ← databe-129..0
        endcase
        StoreMemory(ba, la, 128 - bs, order, data0)
        StoreMemory(ba, (la63..5 + 1) || 04, be - 128, order, data1)
    else
        for i ← 0 to size-8 by 8
            j ← bs + ((order=L) ? i : size-8-i)
            hdataj+7..j ← datai+7..i
        endfor
        xdata ← TranslateAndCacheAccess(ba, la, size, W, hdata)
    endif
endif

```

```

        endif
enddef

```

Memory

Memory operations require first translating via the LTB and GTB, checking for access exceptions, then accessing the cache.

Definition

```

def hdata ← TranslateAndCacheAccess(ba,la,size,rwxg,hwdata)
    if ControlRegister62 then
        case rwxg of
            R:
                at ← 0
            W:
                at ← 1
            X:
                at ← 2
            G:
                at ← 3
        endcase
        rw ← (rwxg=W) ? W : R
        ga,LocalProtect ← LocalTranslation(th,ba,la,pl)
        if LocalProtect9+2*at..8+2*at < pl then
            raise AccessDisallowedByLTB
        endif
        lda ← LocalProtect4
        pa,GlobalProtect ← GlobalTranslation(th,ga,pl,lda)
        if GlobalProtect9+2*at..8+2*at < pl then
            raise AccessDisallowedByGTB
        endif
        cc ← (LocalProtect2..0 > GlobalProtect2..0) ? LocalProtect2..0 : GlobalProtect2..0
        so ← LocalProtect3 or GlobalProtect3
        gda ← GlobalProtect4
        hdata,TagProtect ← LevelOneCacheAccess(pa,size,lda,gda,cc,rw,hwdata)
        if (lda ^ gda ^ TagProtect) = 1 then
            if TagProtect then
                PerformAccessDetail(AccessDetailRequiredByTag)
            elseif gda then
                PerformAccessDetail(AccessDetailRequiredByGlobalTB)
            else
                PerformAccessDetail(AccessDetailRequiredByLocalTB)
            endif
        endif
    else
        case rwxg of
            R, X, G:
                hdata ← ReadPhysical(la,size)
            W:
                WritePhysical(la,size,hwdata)
        endcase
    endif
enddef

```


Bus interface

The initial implementation of the Zeus processor uses a “Super Socket 7 compatible” (SS7) bus interface, which is generally similar to and compatible with other “Socket 7” and “Super Socket 7” processors such as the Intel Pentium, Pentium with MMX Technology; AMD K6, K6-II, K6-III; IDT Winchip C6, 2, 2A, 3, 4; Cyrix 6x86, etc. and other “Socket 7” chipsets listed below.

The SS7 bus interface behavior is quite complex, but well-known due to the leading position of the Intel Pentium design. This document does not yet contain all the detailed information related to this bus, and will concentrate on the differences between the Zeus SS7 bus and other designs. For functional specification and pin interface behavior, the *Pentium Processor Family Developer's Manual²⁷* is a primary reference. For 100 MHz SS7 bus timing data, the *AMD K6-2 Processor Data Sheet²⁸* is a primary reference.

Motherboard Chipsets

The following motherboard chipsets are designed for the 100 MHz “Socket 7” bus:

Manufacturer	Website	Chipset	clock rate	North bridge	South bridge
VIA technologies, Inc.	www_via_com_tw	Apollo MVP3	100 MHz	vt82c598at ²⁹	vt82c598b
Silicon Integrated Systems	www_sis_com_tw	SIS 5591/5592	75 MHz	SIS 5591 ³⁰	SIS 5595
Acer Laboratories, Inc.	www_acerlabs_com	Ali Aladdin V	100 MHz	M1541 ³¹	M1543C

The following processors are designed for a “Socket 7” bus:

Manufacturer	Website	Chips	clock rate
Advanced Micro Devices	www_amd_com	K6-2	100 MHz
Advanced Micro Devices	www_amd_com	K6-3	100 MHz
Intel	www_intel_com	Pentium MMX	66 MHz
IDT/Centaur	www_winchip_com	Winchip C6	75 MHz
IDT/Centaur	www_winchip_com	Winchip 2	100 MHz
IDT/Centaur	www_winchip_com	Winchip 2A	100 MHz
IDT/Centaur	www_winchip_com	Winchip 4	100 MHz
NSM/Cyrix	www_cyrix_com		

²⁷ <http://home.microunity.com/~craig/standards/intel/intel-pentium-family-developers-manual-1997-24142805.pdf>

²⁸ <http://home.microunity.com/~craig/standards/amd/amd-k62-data-sheet-21850c.pdf>

²⁹ <http://home.microunity.com/~craig/standards/via/via-apollo-mvp3-vt82c598at-598.pdf>

³⁰ <http://home.microunity.com/~craig/standards/sis/5591ds10.doc>

³¹ <http://home.microunity.com/~craig/standards/acer/aladdin5pb.htm>

Pinout

In the diagram below, signals which are different from Pentium pinout, are indicated by italics and underlining. Generally, other Pentium-compatible processors (such as the AMD K6-2) define these signals.

	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1					
AN	VSS	NC	A6	A10	VCC3	VCC3	VCC3	VCC3	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	FLUSH#	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	INC	AN				
AM	A30	A4	A8	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	W/R#	EADS#	ADSC#	AM	AM	AM	AM	AM	AM															
AL	VSS	A3	A7	A11	A12	A14	A16	A18	A20	NC	SCYC	BE6#	BE4#	BE2#	BE0#	BUSCK#	HITM#	PWT	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	AL				
AK	A28	A29	A5	A9	A13	A15	A17	A19	RESET	CLK	BE7#	BE5#	BE3#	BE1#	A20M#	HIT#	D/C#	AP	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK	AK			
AJ	VSS	A25	A31																																		AJ					
AH		A22	A26	<u>KEY</u>																																	AH					
AG	VCC3	A24	A27																																		AG					
AF	VSS	A21																																			AF					
AE	VCC3	D/P#	A23																																		AE					
AD	VSS	INTR																																			AD					
AC	VCC3	R/S#	NMI																																		AC					
AB	VSS	SM#																																			AB					
AA	VCC3	IGNNE#	INIT																																		AA					
Z	VSS	PEN#																																			Z					
Y	VCC3	FRCMC#	BFO																																		Y					
X	VSS	BF1																																			X					
W	VCC3	<u>BF2</u>	NC																																		W					
V	VSS	STPCLK#																																			V					
U	VCC3	VSS	VCC3																																		U					
T	VSS	VCC3																																			T					
S	VCC3	NC	NC																																		S					
R	VSS	NC																																			R					
Q	VCC3	CPUTP	TRST#																																		Q					
P	VSS	TMS																																			P					
N	VCC3	TDI	TDO																																		N					
M	VSS	TCK																																			M					
L	VCC3	PICD1	VCC3																																		L					
K	VSS	DO																																			K					
J	VCC3	D2	PICD0																																		J					
H	VSS	PICCLK																																			H					
G	VCC3	D1	D3																																		G					
F	D4	D5																																			F					
E	VSS	A25	A31																																		E					
D	DP0	D8	D12	DP1	D19	D23	D26	D28	D30	DP3	D33	D35	D37	D39	D40	D44	D48	D58																			D					
C	D9	D10	D14	D17	D21	D24	DP2	D25	D27	D29	D31	D32	D34	D36	D38	DP4	D45	D47	INC																			C				
B	D11	D13	D16	D20	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	D43	INC																		B			
A	NC	D15	D18	D22	VCC3	VCC3	VCC3	VCC3	VCC3	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	VCC2	D41	INC																		A			

Pin summary

A20M#	I	Address bit 20 Mask is an emulator signal.
A31..A3	IO	Address, in combination with byte enable, indicate the physical addresses of memory or device that is the target of a bus transaction. This signal is an output,

		when the processor is initiating the bus transaction, and an input when the processor is receiving an inquire transaction or snooping another processor's bus transaction.
ADS#	IO	ADdress Strobe , when asserted, indicates new bus transaction by the processor, with valid address and byte enable simultaneously driven.
ADSC#	O	Address Strobe Copy is driven identically to address strobe
AHOLD	I	Address HOLD , when asserted, causes the processor to cease driving address and address parity in the next bus clock cycle.
AP	IO	Address Parity contains even parity on the same cycle as address . Address parity is generated by the processor when address is an output, and is checked when address is an input. A parity error causes a bus error machine check.
APCHK#	O	Address Parity CHecK is asserted two bus clocks after EADS# if address parity is not even parity of address .
APICEN	I	Advanced Programmable Interrupt Controller ENable is not implemented.
BE7#..BE0#	IO	Byte Enable indicates which bytes are the subject of a read or write transaction and are driven on the same cycle as address .
BF1..BF0	I	Bus Frequency is sampled to permit software to select the ratio of the processor clock to the bus clock.
BOFF#	I	BackOFF is sampled on the rising edge of each bus clock, and when asserted, the processor floats bus signals on the next bus clock and aborts the current bus cycle, until the backoff signal is sampled negated.
BP3..BPO	O	BreakPoint is an emulator signal.
BRDY#	I	Bus ReaDY indicates that valid data is present on data on a read transaction, or that data has been accepted on a write transaction.
BRDYC#	I	Bus ReaDY Copy is identical to BRDY#; asserting either signal has the same effect.
BREQ	O	Bus REQuest indicates a processor initiated bus request.
BUSCHK#	I	BUS CHecK is sampled on the rising edge of the bus clock, and when asserted, causes a bus error machine check.
CACHE#	O	CACHE , when asserted, indicates a cacheable read transaction or a burst write transaction.
CLK	I	bus CLock provides the bus clock timing edge and the frequency reference for the processor clock.
CPUTYP	I	CPU TYP , if low indicates the primary processor, if high, the dual processor.

D/C#	I	Data/Cod is driven with the address signal to indicate data, code, or special cycles.
D63..D0	IO	Data communicates 64 bits of data per bus clock .
D/P#	O	Dual/Primary is driven (asserted, low) with address on the primary processor
DP7..DP0	IO	Data Parity contains even parity on the same cycle as data . A parity error causes a bus error machine check.
DPEN#	IO	Dual Processing Enable is asserted (driven low) by a Dual processor at reset and sampled by a Primary processor at the falling edge of reset.
EADS#	I	External Address Strobe indicates that an external device has driven address for an inquire cycle.
EWBE#	I	External Write Buffer Empty indicates that the external system has no pending write.
FERR#	O	Floating point ERROR is an emulator signal.
FLUSH#	I	cache FLUSH is an emulator signal.
FRCMC#	I	Functional Redundancy Checking Master/Checker is not implemented.
HIT#	IO	HIT indicates that an inquire cycle or cache snoop hits a valid line.
HITM#	IO	HIT to a Modified Line indicates that an inquire cycle or cache snoop hits a sub-block in the M cache state.
HLDA	O	bus HOLD Acknowledge is asserted (driven high) to acknowledge a bus hold request
HOLD	I	bus HOLD request causes the processor to float most of its pins and assert bus hold acknowledge after completing all outstanding bus transactions, or during reset.
IERR#	O	Internal ERROR is an emulator signal.
IGNNE#	I	IGNore Numeric Error is an emulator signal.
INIT	I	INITialization is an emulator signal.
INTR	I	maskable INTerrupt is an emulator signal.
INV	I	INVALIDation controls whether to invalidate the addressed cache sub-block on an inquire transaction.
KEN#	I	Cache ENable is driven with address to indicate that the read or write transaction is cacheable.
LINT1..LINT0	I	Local INTerrupt is not implemented.
LOCK#	O	bus LOCK is driven starting with address and ending after bus ready to indicate a locked series of bus transactions.
M/IO#	O	Memory/Input Output is driven with address to indicate a memory or I/O transaction.
NA#	I	Next Address indicates that the external system will accept an address for a new bus cycle in two bus clocks.
NMI	I	Non Maskabl Interrupt is an emulator signal.

PBGNT#	IO	Private Bus GraNT is driven between Primary and Dual processors to indicate that bus arbitration has completed, granting a new master access to the bus.
PBREQ#	IO	Private Bus REQuest is driven between Primary and Dual processors to request a new master access to the bus.
PCD	O	Page Cache Disable is driven with address to indicate a not cacheable transaction.
PCHK#	O	Parity CHecK is asserted (driven low) two bus clocks after data appears with odd parity on enabled bytes.
PHIT#	IO	Private HIT is driven between Primary and Dual processors to indicate that the current read or write transaction addresses a valid cache sub-block in the slave processor.
PHITM#	IO	Private HIT Modified is driven between Primary and Dual processors to indicate that the current read or write transaction addresses a modified cache sub-block in the slave processor.
PICCLK	I	Programmable Interrupt Controller CLock is not implemented.
PICD1..PICD0	IO	Programmable Interrupt Controller Data is not implemented.
PEN#	I	Parity Enable , if active on the data cycle, allows a parity error to cause a bus error machine check.
PM1..PM0	O	Performance Monitoring is an emulator signal.
PRDY	O	Probe ReaDY is not implemented.
PWT	O	Page Write Through is driven with address to indicate a not write allocate transaction.
R/S#	I	Run/Stop is not implemented.
RESET	I	RESET causes a processor reset.
SCYC	O	Split CYCle is asserted during bus lock to indicate that more than two transactions are in the series of bus transactions.
SMI#	I	System Management Interrupt is an emulator signal.
SMIACT#	O	System Management Interrupt ACTive is an emulator signal.
STPCLK#	I	SToP CLock is an emulator signal.
TCK	I	Test CLock follows IEEE 1149.1.
TDI	I	Test Data Input follows IEEE 1149.1.
TDO	O	Test Data Output follows IEEE 1149.1.
TMS	I	Test Mode Select follows IEEE 1149.1.
TRST#	I	Test ReSeT follows IEEE 1149.1.
VCC2	I	VCC of 2.8V at 25 pins
VCC3	I	VCC of 3.3V at 28 pins
VCC2DET#	O	VCC2 DETect sets appropriate VCC2 voltage level.
VSS	I	VSS supplied at 53 pins
W/R#	O	Write/Read is driven with address to indicate write

		vs. read transaction.
WB/WT#	I	Write Back/Write Through is returned to indicate that data is permitted to be cached as write back.

Electrical Specifications

These preliminary electrical specifications provide AC and DC parameters that are required for "Super Socket 7" compatibility.

Clock rate	66 MHz		75 MHz		100 MHz		133 MHz		
Parameter	min	max	min	max	min	max	min	max	unit
CLK frequency	33.3	66.7	37.5	75	50	100		133	MHz
CLK period	15.0	30.0	13.3	26.3	10.0	20.0			ns
CLK high time ($\geq 2V$)	4.0		4.0		3.0				ns
CLK low time ($\leq 0.8V$)	4.0		4.0		3.0				ns
CLK rise time (0.8V->2V)	0.15	1.5	0.15	1.5	0.15	1.5			ns
CLK fall time (2V->0.8V)	0.15	1.5	0.15	1.5	0.15	1.5			ns
CLK period stability		250		250		250			ps

A31..3 valid delay	1.1	6.3	1.1	4.5	1.1	4.0			ns
A31..3 float delay		10.0		7.0		7.0			ns
ADS# valid delay	1.0	6.0	1.0	4.5	1.0	4.0			ns
ADS# float delay		10.0		7.0		7.0			ns
ADSC# valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
ADSC# float delay		10.0		7.0		7.0			ns
AP valid delay	1.0	8.5	1.0	5.5	1.0	5.5			ns
AP float delay		10.0		7.0		7.0			ns
APCHK# valid delay	1.0	8.3	1.0	4.5	1.0	4.5			ns
BE7..0# valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
BE7..0# float delay		10.0		7.0		7.0			ns
BP3..0 valid delay	1.0	10.0							ns
BREQ valid delay	1.0	8.0	1.0	4.5	1.0	4.0			ns
CACHE# valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
CACHE# float delay		10.0		7.0		7.0			ns
D/C# valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
D/C# float delay		10.0		7.0		7.0			ns
D63..0 write data valid delay	1.3	7.5	1.3	4.5	1.3	4.5			ns
D63..0 write data float delay		10.0		7.0		7.0			ns
DP7..0 write data valid delay	1.3	7.5	1.3	4.5	1.3	4.5			ns
DP7..0 write data float delay		10.0		7.0		7.0			ns
FERR# valid delay	1.0	8.3	1.0	4.5	1.0	4.5			ns
HIT# valid delay	1.0	6.8	1.0	4.5	1.0	4.0			ns
HITM# valid delay	1.1	6.0	1.1	4.5	1.1	4.0			ns
HLDA valid delay	1.0	6.8	1.0	4.5	1.0	4.0			ns
IERR# valid delay	1.0	8.3							ns
LOCK# valid delay	1.1	7.0	1.1	4.5	1.1	4.0			ns
LOCK# float delay		10.0		7.0		7.0			ns
M/IO# valid delay	1.0	5.9	1.0	4.5	1.0	4.0			ns
M/IO# float delay		10.0		7.0		7.0			ns
PCD valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
PCD float delay		10.0		7.0		7.0			ns
PCHK# valid delay	1.0	7.0	1.0	4.5	1.0	4.5			ns
PM1..0 valid delay	1.0	10.0							ns
PRDY valid delay	1.0	8.0							ns
PWT valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
PWT float delay		10.0		7.0		7.0			ns
SCYC valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
SCYC float delay		10.0		7.0		7.0			ns
SMIACT# valid delay	1.0	7.3	1.0	4.5	1.0	4.0			ns
W/R# valid delay	1.0	7.0	1.0	4.5	1.0	4.0			ns
W/R# float delay		10.0		7.0		7.0			ns

A31..5 setup time	6.0	3.0	3.0			ns
A31..5 hold time	1.0	1.0	1.0			ns
A20M# setup time	5.0	3.0	3.0			ns
A20M# hold time	1.0	1.0	1.0			ns
AHOLD setup time	5.5	3.5	3.5			ns
AHOLD hold time	1.0	1.0	1.0			ns
AP setup time	5.0	1.7	1.7			ns
AP hold time	1.0	1.0	1.0			ns
BOFF# setup time	5.5	3.5	3.5			ns
BOFF# hold time	1.0	1.0	1.0			ns
BRDY# setup time	5.0	3.0	3.0			ns
BRDY# hold time	1.0	1.0	1.0			ns
BRDYC# setup time	5.0	3.0	3.0			ns
BRDYC# hold time	1.0	1.0	1.0			ns
BUSCHK# setup time	5.0	3.0	3.0			ns
BUSCHK# hold time	1.0	1.0	1.0			ns
D63..0 read data setup time	2.8	1.7	1.7			ns
D63..0 read data hold time	1.5	1.5	1.5			ns
DP7..0 read data setup time	2.8	1.7	1.7			ns
DP7..0 read data hold time	1.5	1.5	1.5			ns
EADS# setup time	5.0	3.0	3.0			ns
EADS# hold time	1.0	1.0	1.0			ns
EWBE# setup time	5.0	1.7	1.7			ns
EWBE# hold time	1.0	1.0	1.0			ns
FLUSH# setup time	5.0	1.7	1.7			ns
FLUSH# hold time	1.0	1.0	1.0			ns
FLUSH# async pulse width	2	2	2			CLK
HOLD setup time	5.0	1.7	1.7			ns
HOLD hold time	1.5	1.5	1.5			ns
IGNNE# setup time	5.0	1.7	1.7			ns
IGNNE# hold time	1.0	1.0	1.0			ns
IGNNE# async pulse width	2	2	2			CLK
INIT setup time	5.0	1.7	1.7			ns
INIT hold time	1.0	1.0	1.0			ns
INIT async pulse width	2	2	2			CLK
INTR setup time	5.0	1.7	1.7			ns
INTR hold time	1.0	1.0	1.0			ns
INV setup time	5.0	1.7	1.7			ns
INV hold time	1.0	1.0	1.0			ns
KEN# setup time	5.0	3.0	3.0			ns
KEN# hold time	1.0	1.0	1.0			ns
NA# setup time	4.5	1.7	1.7			ns
NA# hold time	1.0	1.0	1.0			ns
NMI setup time	5.0	1.7	1.7			ns
NMI hold time	1.0	1.0	1.0			ns
NMI async pulse width	2	2	2			CLK

<u>PEN# setup time</u>	4.8	1.7	1.7			ns
<u>PEN# hold time</u>	1.0	1.0	1.0			ns
<u>R/S# setup time</u>	5.0	1.7	1.7			ns
<u>R/S# hold time</u>	1.0	1.0	1.0			ns
<u>R/S# async pulse width</u>	2	2	2			CLK
<u>SMI# setup time</u>	5.0	1.7	1.7			ns
<u>SMI# hold time</u>	1.0	1.0	1.0			ns
<u>SMI# async pulse width</u>	2	2	2			CLK
<u>STPCLK# setup time</u>	5.0	1.7	1.7			ns
<u>STPCLK# hold time</u>	1.0	1.0	1.0			ns
<u>WB/WT# setup time</u>	4.5	1.7	1.7			ns
<u>WB/WT# hold time</u>	1.0	1.0	1.0			ns

<u>RESET setup time</u>	5.0	1.7	1.7			ns
<u>RESET hold time</u>	1.0	1.0	1.0			ns
<u>RESET pulse width</u>	15	15	15			CLK
<u>RESET active</u>	1.0	1.0	1.0			ms
<u>BF2..0 setup time</u>	1.0	1.0	1.0			ms
<u>BF2..0 hold time</u>	2	2	2			CLK
<u>BRDYC# hold time</u>	1.0	1.0	1.0			ns
<u>BRDYC# setup time</u>	2	2	2			CLK
<u>BRDYC# hold time</u>	2	2	2			CLK
<u>FLUSH# setup time</u>	5.0	1.7	1.7			ns
<u>FLUSH# hold time</u>	1.0	1.0	1.0			ns
<u>FLUSH# setup time</u>	2	2	2			CLK
<u>FLUSH# hold time</u>	2	2	2			CLK

PBREQ# flight time	0	2.0						ns
PBGNT# flight time	0	2.0						ns
PHIT# flight time	0	2.0						ns
PHITM# flight time	0	1.8						ns
A31..5 setup time	3.7							ns
A31..5 hold time	0.8							ns
D/C# setup time	4.0							ns
D/C# hold time	0.8							ns
W/R# setup time	4.0							ns
W/R# hold time	0.8							ns
CACHE# setup time	4.0							ns
CACHE# hold time	1.0							ns
LOCK# setup time	4.0							ns
LOCK# hold time	0.8							ns
SCYC setup time	4.0							ns
SCYC hold time	0.8							ns
ADS# setup time	5.8							ns
ADS# hold time	0.8							ns
M/IO# setup time	5.8							ns
M/IO# hold time	0.8							ns
HIT# setup time	6.0							ns
HIT# hold time	1.0							ns
HITM# setup time	6.0							ns
HITM# hold time	0.7							ns
HLDA setup time	6.0							ns
HLDA hold time	0.8							ns
DPEN# valid time		10.0						CLK
DPEN# hold time	2.0							CLK
D/P# valid delay (primary)	1.0	8.0						ns

TCK frequency		25			25			MHz
TCK period	40.0				40.0			ns
TCK high time ($\geq 2V$)	14.0				14.0			ns
TCK low time ($\leq 0.8V$)	14.0				14.0			ns
TCK rise time (0.8V->2V)		5.0				5.0		ns
TCK fall time (2V->0.8V)		5.0				5.0		ns
TRST# pulse width	30.0				30.0			ns

TDI setup time	5.0			5.0			ns
TDI hold time	9.0			9.0			ns
TMS setup time	5.0			5.0			ns
TMS hold time	9.0			9.0			ns
TDO valid delay	3.0	13.0		3.0	13.0		ns
TDO float delay		16.0			16.0		ns
all outputs valid delay	3.0	13.0		3.0	13.0		ns
all outputs float delay		16.0			16.0		ns
all inputs setup time	5.0			5.0			ns
all inputs hold time	9.0			9.0			ns

Bus Control Register

The Bus Control Register provides direct control of Emulator signals, selecting output states and active input states for these signals.

The layout of the Bus Control Register is designed to match the assignment of signals to the Event Register.

number	control
0	Reserved
1	A20M# active level
2	BF0 active level
3	BF1 active level
4	BF2 active level
5	BUSCHK active level
6	FLUSH# active level
7	FRCMC# active level
8	IGNNE# active level
9	INIT active level
10	INTR active level
11	NMI active level
12	SMI# active level
13	STPCLK# active level
14	CPUTYP active at reset
15	DPEN# active at reset
16	FLUSH# active at reset
17	INIT active at reset
31..18	Reserved
32	Bus lock
33	Split cycle
34	BP0 output
35	BP1 output
36	BP2 output
37	BP3 output
38	FERR# output
39	IERR# output
40	PM0 output
41	PM1 output
42	SMIACT# output
63..43	Reserved

Emulator signals

Several of the signals, A20M#, INIT, NMI, SMI#, STPCLK#, IGNNE# are inputs that have purposes primarily defined by the needs of x86 processor emulation. They have no direct purpose in the Zeus processor, other than to signal an event, which is handled by software. Each of these signals is an input sampled on the rising edge of each bus clock, if the input signal matches the active level specified in the **bus control register**, the corresponding bit in the **event register** is set. The bit in the **event register** remains set even if the signal is no longer active, until cleared by software. If the **event register** bit is cleared by software, it is set again on each **bus clock** that the signal is sampled active.

A20M#

A20M# (address bit 20 mask inverted), when asserted (low), directs an x86 emulator to generate physical addresses for which bit 20 is zero.

The **A20M#** bit of the **bus control register** selects which level of the **A20M#** signal will generate an event in the **A20M#** bit of the **event register**. Clearing (to 0) the **A20M#** bit of the **bus control register** will cause the **A20M#** bit of the **event register** to be set when the **A20M#** signal is asserted (low).

Asserting the **A20M#** signal causes the emulator to modify all current TB mappings to produce a zero value for bit 20 of the byte address. The **A20M#** bit of the **bus control register** is then set (to 1) to cause the **A20M#** bit of the **event register** to be set when the **A20M#** signal is released (high).

Releasing the **A20M#** signal causes the emulator to restore the TB mapping to the original state. The **A20M#** bit of the **bus control register** is then cleared (to 0) again, to cause the **A20M#** bit of the **event register** to be set when the **A20M#** signal is asserted (low).

INIT

INIT (initialize) when asserted (high), directs an x86 emulator to begin execution of the external ROM BIOS.

The **INIT** bit of the **bus control register** is normally set (to 1) to cause the **INIT** bit of the **event register** to be set when the **INIT** signal is asserted (high).

INTR

INTR (maskable interrupt) when asserted (high), directs an x86 emulator to simulate a maskable interrupt by generating two locked **interrupt acknowledge** special cycles. External hardware will normally release the **INTR** signal between the first and second **interrupt acknowledge** special cycle.

The **INTR** bit of the **bus control register** is normally set (to 1) to cause the **INTR** bit of the **event register** to be set when the **INTR** signal is asserted (high).

NMI

NMI (non-maskable interrupt) when asserted (high), directs an x86 emulator to simulate a non-maskable interrupt. External hardware will normally release the **NMI** signal.

The **NMI** bit of the **bus control register** is normally set (to 1) to cause the **NMI** bit of the **event register** to be set when the **NMI** signal is asserted (high).

SMI#

SMI# (system management interrupt inverted) when asserted (low), directs an x86 emulator to simulate a system management interrupt by flushing caches and saving registers, and asserting (low) **SMIACT#** (system management interrupt active inverted). External hardware will normally release the **SMI#**.

The **SMI#** bit of the **bus control register** is normally cleared (to 0) to cause the **SMI#** bit of the **event register** to be set when the **SMI#** signal is asserted (low).

STPCLK#

STPCLK# (stop clock inverted) when asserted (low), directs an x86 emulator to simulate a stop clock interrupt by flushing caches and saving registers, and performing a **stop grant** special cycle.

The **STPCLK#** bit of the **bus control register** is normally cleared (to 0) to cause the **STPCLK#** bit of the **event register** to be set when the **STPCLK#** signal is asserted (low).

Software must set (to 1) the **STPCLK#** bit of the **bus control register** to cause the **STPCLK#** bit of the **event register** to be set when the **STPCLK#** signal is released (high) to resume execution. Software must cease producing bus operations after the **stop grant** special cycle. Usually, software will use the **B.HALT** instruction in all threads to cease performing operations. The processor PLL continues to operate, and the processor must still sample INIT, INTR, RESET, NMI, SMI# (to place them in the event register) and respond to RESET and inquire and snoop transactions, so long as the **bus clock** continues operating.

The **bus clock** itself cannot be stopped until the **stop grant** special cycle. If the **bus clock** is stopped, it must stop in the low (0) state. The **bus clock** must be operating at frequency for at least 1 ms before releasing **STPCLK#** or releasing **RESET**. While the **bus clock** is stopped, the processor does not sample inputs or responds to **RESET** or **inquire** or **snoop** transactions.

External hardware will normally release **STPCLK#** when it is desired to resume execution. The processor should respond to the **STPCLK#** bit in the **event register** by awakening one or more threads.

IGNNE#

IGNNE# (address bit 20 mask inverted), when asserted (low), directs an x86 emulator to ignore numeric errors.

The **IGNNE#** bit of the **bus control register** selects which level of the **IGNNE#** signal will generate an event in the **IGNNE#** bit of the **event register**. Clearing (to 0) the **IGNNE#** bit of the **bus control register** will cause the **IGNNE#** bit of the **event register** to be set when the **IGNNE#** signal is asserted (low).

Asserting the IGNNE# signal causes the emulator to modify its processing to ignore numeric errors, if suitably enabled to do so. The IGNNE# bit of the bus control register is then set (to 1) to cause the IGNNE# bit of the event register to be set when the IGNNE# signal is released (high).

Releasing the IGNNE# signal causes the emulator to restore the emulation to the original state. The IGNNE# bit of the bus control register is then cleared (to 0) again, to cause the IGNNE# bit of the event register to be set when the IGNNE# signal is asserted (low).

Emulator output signals

Several of the signals, BP3..BP0, FERR#, IERR#, PM1..PM0, SMIACT# are outputs that have purposes primarily defined by the needs of x86 processor emulation. They are driven from the bus control register that can be written by software.

Bus snooping

Zeus support the "Socket 7" protocols for inquiry, invalidation and coherence of cache lines. The protocols are implemented in hardware and do not interrupt the processor as a result of bus activity. Cache access cycles may be "stolen" for this purpose, which may delay completion of processor memory activity.

Definition

```
def SnoopPhysicaBus as
    //wait for transaction on bus or inquiry cycle
    do
        wait
    while BRDY# = 0
    pa31..3 <- A31..3
    op <- W/R# ? W : R
    cc <- CACHE# || PWT || PCD
enddef
```

Locked cycles

Locked cycles occur as a result of synchronization operations (Store-swap instructions) performed by the processor. For x86 emulation, locked cycles also occur as a result of setting specific memory-mapped control registers.

Locked synchronization instruction

Bus lock (LOCK#) is asserted (low) automatically as a result of store-swap instructions that generate bus activity, which always perform locked read-modify-write cycles on 64 bits of data. Note that store-swap instructions that are performed on cache sub-blocks that are in the E or M state need not generate bus activity.

Locked sequences of bus transactions

Bus lock (LOCK#) is also asserted (low) on subsequent bus transactions by writing a one (1) to the bus lock bit of the bus control register. Split cycle (SCYC) is similarly asserted (high) if a one (1) is also written to the split cycle bit of the bus emulation control register.

All subsequent bus transactions will be performed as a locked sequence of transactions, asserting bus lock (LOCK# low) and optionally split cycle (SCYC high), until zeroes (0) are written to the bus lock and split cycle bits of the bus control register. The next bus transaction completes the locked sequence, releasing bus lock (LOCK# high) and split cycle (SCYC low) at the end of the transaction. If the locked transaction must be aborted because of bus activity such as backoff, a lock broken event is signalled and the bus lock is released.

Unless special care is taken, the bus transactions of all threads occur as part of the locked sequence of transactions. Software can do so by interrupting all other threads until the locked sequence is completed. Software should also take care to avoid fetching instructions during the locked sequence, such as by executing instructions out of niche or ROM memory. Software should also take care to avoid terminating the sequence with event handling prior to releasing the bus lock, such as by executing the sequence with events disabled (other than the lock broken event).

The purpose of this facility is primarily for x86 emulation purposes, in which we are willing to perform vile acts (such as stopping all the other threads) in the name of compatibility. It is possible to take special care in hardware to sort out the activity of other threads, and break the lock in response to events. In doing so, the bus unit must defer bus activity generated by other threads until the locked sequence is completed. The bus unit should inhibit event handling while the bus is locked.

Sampled at Reset

Certain pins are sampled at reset and made available in the event register.

CPUTYP Primary or Dual processor

PICD0[DPEN#] Dual processing enable

FLUSH# Tristate test mode

INIT Built-in self-test

Sampled per Clock

Certain pins are sampled per clock and changes are made available in the event register.

A20M# address bit 20 mask

BF[1:0] bus frequency

BUSCHK# bus check

FLUSH# cache flush request

FRCMC#functional redundancy check - not implemented on Pentium MMX

IGNNE#ignore numeric error

INIT re-initialize pentium processor

INTR external interrupt

NMI non-maskable interrupt

R/S# run/stop

SMI# system management

STPCLK# stop clock

Bus Access

The "Socket 7" bus performs transfers of 1-8 bytes within an octlet boundary or 32 bytes on a triclet boundary.

Transfers sized at 16 bytes (hexlet) are not available as a single transaction, they are performed as two bus transactions.

Bus transactions begin by gaining control of the bus (TODO: not shown), and in the initial cycle, asserting ADS#, M/IO#, A, BE#, W/R#, CACHE#, PWT, and PCD. These signals indicate the type, size, and address of the transaction. One or more octlets of data are returned on a read (the external system asserts BRDY# and/or NA# and D), or accepted on a write (TODO not shown).

The external system is permitted to affect the cacheability and exclusivity of data returned to the processor, using the KEN# and WB/WT# signals.

Definition

```
def data,cen ← AccessPhysicalBus(pa,size,cc,op,wd) as
    // divide transfers sized between octlet and hexlet into two parts
    // also divide transfers which cross octlet boundary into two parts
    if {64<size≤128} or ({size<64} and {size+8*pa2..0>64}) then
        data0,cen ← AccessPhysicalBus(pa,64-8*pa2..0,cc,op,wd)
        if cen=0 then
            pa1 ← pa63..41111103
            data1,cen ← AccessPhysicalBus(pa1,size+8*pa2..0-64,cc,op,wd)
            data ← data1127..64 11 data063..0
        endif
    endif
```

```

else
    ADS# ← 0
    M/IO# ← 1
    A31..3 ← pa31..3
    for i ← 0 to 7
        BEi# ← pa2..0 ≤ i < pa2..0+size/8
    endfor
    W/R# ← {op = W}
    if {op=R} then
        CACHE# ← ~{cc ≥ WT}
        PWT ← {cc = WT}
        PCD ← {cc ≤ CD}
        do
            wait
            while {BRDY# = 1} and {NA# = 1}
                //Intel spec doesn't say whether KEN# should be ignored if no CACHE#
                //AMD spec says KEN# should be ignored if no CACHE#
                cen ← ~KEN# and {cc ≥ WT} //cen=1 if triclet is cacheable
                xen ← WB/WT# and {cc ≠ WT} //xen=1 if triclet is exclusive
            if cen then
                os ← 64*pa4..3
                data63+os..os ← D63..0
                do
                    wait
                    while BRDY# = 1
                        data63+(64^os)..(64^os) ← D63..0
                    do
                        wait
                        while BRDY# = 1
                            data63+(128^os)..(128^os) ← D63..0
                    do
                        wait
                        while BRDY# = 1
                            data63+(192^os)..(192^os) ← D63..0
            else
                os ← 64*pa3
                data63+os..os ← D63..0
            endif
        else
            CACHE# ← ~{size = 256}
            PWT ← {cc = WT}
            PCD ← {cc ≤ CD}
            do
                wait
                while {BRDY# = 1} and {NA# = 1}
                    xen ← WB/WT# and {cc ≠ WT}
            endif
        endif
        flags ← cen || xen
    enddef

```

Other bus cycles

Input/Output transfers, Interrupt acknowledge and special bus cycles (stop grant, flush acknowledge, writeback, halt, flush, shutdown) are performed by uncached loads and stores to a memory-mapped control region.

M/IO#	D/C#	W/R#	CACHE#	KEN#	cycle
0	0	0	1	x	interrupt acknowledge
0	0	1	1	x	special cycles (intel pg 6-33)
0	1	0	1	x	I/O read, 32-bits or less, non-cacheable, 16-bit address
0	1	1	1	x	I/O write, 32-bits or less, non-cacheable, 16-bit address
1	0	x	x	x	code read (not implemented)
1	1	0	1	x	non-cacheable read
1	1	0	x	1	non-cacheable read
1	1	0	0	0	cacheable read
1	1	1	1	x	non-cacheable write
1	1	1	0	x	cache writeback

Special cycles

An interrupt acknowledge cycle is performed by two byte loads to the control space ($dc=1$), the first with a byte address (ba) of 4 (A31..3=0, BE4#=0, BE7..5,3..0#=1), the second with a byte address (ba) of 0 (A31..3=0, BE0#=0, BE7..1#=1). The first byte read is ignored; the second byte contains the interrupt vector. The external system normally releases INTR between the first and second byte load.

A shutdown special cycle is performed by a byte store to the control space ($dc=1$) with a byte address (ba) of 0 (A31..3=0, BE0#=0, BE7..1#=1).

A flush special cycle is performed by a byte store to the control space ($dc=1$) with a byte address (ba) of 1 (A31..3=0, BE1#=0, BE7..2,0#=1).

A halt special cycle is performed by a byte store to the control space ($dc=1$) with a byte address (ba) of 2 (A31..3=0, BE2#=0, BE7..3,1..0#=1).

A stop grant special cycle is performed by a byte store to the control space ($dc=1$) with a byte address (ba) of 0x12 (A31..3=2, BE2#=0, BE7..3,1..0#=1).

A writeback special cycle is performed by a byte store to the control space ($dc=1$) with a byte address (ba) of 3 (A31..3=0, BE3#=0, BE7..4,2..0#=1).

A flush acknowledge special cycle is performed by a byte store to the control space ($dc=1$) with a byte address (ba) of 4 (A31..3=0, BE4#=0, BE7..5,3..0#=1).

A back trace message special cycle is performed by a byte store to the control space (**dc=1**) with a byte address (**ba**) of 5 (A31..3=0, BE5#=0, BE7..6,4..0#=1).

Performing load or store operations of other sizes (doublet, quadlet, octlet, hexlet) to the control space (**dc=1**) or operations with other byte address (**ba**) values produce bus operations which are not defined by the "Super Socket 7" specifications and have undefined effect on the system.

I/O cycles

An input cycle is performed by a byte, doublet, or quadlet load to the data space (**dc=0**), with a byte address (**ba**) of the I/O address. The address may not be aligned, and if it crosses an octlet boundary, will be performed as two separate cycles.

An output cycle is performed by a byte, doublet, or quadlet store to the data space (**dc=0**), with a byte address (**ba**) of the I/O address. The address may not be aligned, and if it crosses an octlet boundary, will be performed as two separate cycles.

Performing load or store operations of other sizes (octlet, hexlet) to the data space (**dc=0**) produce bus operations which are not defined by the "Super Socket 7" specifications and have undefined effect on the system.

Physical address

The other bus cycles are accessed explicitly by uncached memory accesses to particular physical address ranges. Appropriately sized load and store operations must be used to perform the specific bus cycles required for proper operations. The **dc** field must equal 0 for I/O operations, and must equal 1 for control operations. Within this address range, bus transactions are sized no greater than 4 bytes (quadlet) and do not cross quadlet boundaries.

The physical address of a other bus cycle data/control **dc**, byte address **ba** is:

63	2423	1615	0
FFFF FFFF 0B00 0000	63..24	dc	ba
40	8	16	

Definition

```
def data ← AccessPhysicaOtherBus(pa,size,op,wd) as
    // divide transfers sized between octlet and hexlet into two parts
    // also divide transfers which cross octlet boundary into two parts
    if (64<size≤128) or ((size<64) and (size+8*pa2..0>64)) then
        data0 ← AccessPhysicaOtherBus(pa,64-8*pa2..0,op,wd)
        pa1 ← pa63..4111103
        data1 ← AccessPhysicaOtherBus(pa1,size+8*pa2..0-64,op,wd)
        data ← data1127..64 || data063..0
    else
        ADS# ← 0
        M/IO# ← 0
```

```
D/C# ← ~pa16
A31..3 ← 016 || pa15..3
for i ← 0 to 7
    BEi# ← pa2..0 ≤ i < pa2..0+size/8
endfor
W/R# ← {op = W}
CACHE# ← 1
PWT ← 1
PCD ← 1
do
    wait
while (BRDY# = 1) and (NA# = 1)
if {op=R} then
    os ← 64*pa3
    data63+os..os ← D63..0
endif
endif
enddef
```

Events and Threads

Exceptions signal several kinds of events: (1) events that are indicative of failure of the software or hardware, such as arithmetic overflow or parity error, (2) events that are hidden from the virtual process model, such as translation buffer misses, (3) events that infrequently occur, but may require corrective action, such as floating-point underflow. In addition, there are (4) external events that cause scheduling of a computational process, such as clock events or completion of a disk transfer.

Each of these types of events require the interruption of the current flow of execution, handling of the exception or event, and in some cases, descheduling of the current task and rescheduling of another. The Zeus processor provides a mechanism that is based on the multi-threaded execution model of Mach. Mach divides the well-known UNIX process model into two parts, one called a task, which encompasses the virtual memory space, file and resource state, and the other called a thread, which includes the program counter, stack space, and other register file state. The sum of a Mach task and a Mach thread exactly equals one UNIX process, and the Mach model allows a task to be associated with several threads. On one processor at any one moment in time, at least one task with one thread is running.

In the taxonomy of events described above, the cause of the event may either be synchronous to the currently running thread, generally types 1, 2, and 3, or asynchronous and associated with another task and thread that is not currently running, generally type 4.

For these events, Zeus will suspend the currently running thread in the current task, saving a minimum of registers, and continue execution at a new program counter. The event handler may perform some minimal computation and return, restoring the current threads' registers, or save the remaining registers and switch to a new task or thread context.

Facilities of the exception, memory management, and interface systems are themselves memory mapped, in order to provide for the manipulation of these facilities by high-level language, compiled code. The sole exception is the register file itself, for which standard store and load instructions can save and restore the state.

Definition

```
def Thread(th) as
    forever
        catch exception
            if !(EventRegister and EventMask[th]) ≠ 0 then
                if ExceptionState=0 then
                    raise EventInterrupt
                endif
            endif
            inst ← LoadMemoryX[ProgramCounter,ProgramCounter,32,L]
            Instruction(inst)
        endcatch
        case exception of
            EventInterrupt,
            ReservedInstruction,
            AccessDisallowedByVirtualAddress,
            AccessDisallowedByTag,
```

```

AccessDisallowedByGlobalTB,
AccessDisallowedByLocalTB,
AccessDetailRequiredByTag,
AccessDetailRequiredByGlobalTB,
AccessDetailRequiredByLocalTB,
MissInGlobalTB,
MissInLocalTB,
FixedPointArithmetic,
FloatingPointArithmetic,
GatewayDisallowed:
    case ExceptionState of
        0:
            PerformException(exception)
        1:
            PerformException(SecondException)
        2:
            raise ThirdException
    endcase
TakenBranch:
    ContinuationState ← (ExceptionState=0) ? 0 : ContinuationState
TakenBranchContinue:
    /* nothing */
none, others:
    ProgramCounter ← ProgramCounter + 4
    ContinuationState ← (ExceptionState=0) ? 0 : ContinuationState
endcase
endforever
enddef

```

Definition

```

def PerformException(exception) as
    v ← (exception > 7) ? 7 : exception
    t ← LoadMemory(ExceptionBase,ExceptionBase+Thread*128+64+8*v,64,L)
    if ExceptionState = 0 then
        u ← RegRead[3,128] || RegRead[2,128] || RegRead[1,128] || RegRead[0,128]
        StoreMemory(ExceptionBase,ExceptionBase+Thread*128,512,L,u)
        RegWrite[0,64,ProgramCounter63..2 || PrivilegeLevel
        RegWrite[1,64,ExceptionBase+Thread*128]
        RegWrite[2,64,exception]
        RegWrite[3,64,FailingAddress]
    endif
    PrivilegeLevel ← t1..0
    ProgramCounter ← t63..2 || 02
    case exception of
        AccessDetailRequiredByTag,
        AccessDetailRequiredByGlobalTB,
        AccessDetailRequiredByLocalTB:
            ContinuationState ← ContinuationState + 1
        others:
            /* nothing */
    endcase
    ExceptionState ← ExceptionState + 1
enddef

```

Definition

```
def PerformAccessDetail(exception) as
    if (ContinuationState = 0) or (ExceptionState ≠ 0) then
        raise exception
    else
        ContinuationState ← ContinuationState - 1
    endif
enddef
```



Definition

```
def BranchBack(rd,rc,rb) as
    c ← RegRead(rc, 64)
    if (rd ≠ 0) or (rc ≠ 0) or (rb ≠ 0) then
        raise ReservedInstruction
    endif
    a ← LoadMemory(ExceptionBase,ExceptionBase+Thread*128,128,L)
    if PrivilegeLevel > c1..0 then
        PrivilegeLevel ← c1..0
    endif
    ProgramCounter ← c63..2 || 02
    ExceptionState ← 0
    RegWrite(rd,128,a)
    raise TakenBranchContinue
enddef
```

The following data is stored into memory at the Exception Storage Address

511	384383	256255	128127	0
	RF[3]	RF[2]	RF[1]	RF[0]
128	128	128	128	128
511				384
	RF[3]			
	128			
383				256
	RF[2]			
	128			
255				128
	RF[1]			
	128			
127				0
	RF[0]			
	128			

The following data is loaded from memory at the Exception Vector Address:



The following data replaces the original contents of RF[3..0]:

511	384383	256255	128127	6463	0
	FA	exception	EDP	eps	pc lpl
128	128	128	64	64	
511		448447			384
	0		FA		
64		64			
383		320319	280383	256261	260259 256
	0	0	as	at	code
64		40	8	8	8
255		192191			128
	0		EDP		
64		64			
127		6463			21 0
	0		pc	pl	
64		62			2

at: access type: 0=r, 1=w, 2=x, 3=g

as: access size in bytes

TODO: add size, access type to exception data in pseudocode.

Ephemeral Program State

Ephemeral Program State (EPS) is defined as program state which affects the operation of certain instructions, but which does not need to be saved and restored as part of user state.

Because these bits are not saved and restored, the sizes and values described here are not visible to software. The sizes and values described here were chosen to be convenient for the definitions in this documentation. Any mapping of these values which does not alter the functions described may be used in a conforming implementation. For example, either of the EPS states may be implemented as a thermometer-coded vector, or the ContinuationState field may be represented with specific values for each AccessDetailRequired exception which an instruction execution may encounter.

There are eight bits of EPS:

bit#	Name	Meaning
1..0	ExceptionState	<p>0: Normal processing. Asynchronous events and Synchronous exceptions enabled.</p> <p>1: Event/Exception handling: Synchronous exceptions cause SecondException. Asynchronous events are masked.</p> <p>2: Second exception handling: Synchronous exceptions cause a machine check. Asynchronous events are masked.</p> <p>3: illegal state</p> <p>This field is incremented by handling an event or exception, and cleared by the Branch Back instruction.</p>
7..2	ContinuationState	<p>Continuation state for AccessDetailRequired exceptions. A value of zero enables all exceptions of this kind. The value is increased by one for each AccessDetailRequired exception handled, for which many AccessDetailRequired exceptions are continued past (ignored) on re-execution in normal processing (ex=0). Any other kind of exception, or the completion of an instruction under normal processing causes the continuation state to be reset to zero. State does not need to be saved on context switch.</p>

The ContinuationState bits are ephemeral because if they are cleared as a result of a context switch, the associated exceptions can happen over again. The AccessDetail exception handlers will then set the bits again, as they were before the context switch. In the case where an AccessDetail exception handler must indicate an error, care must be taken to perform some instruction at the target of the Branch Back instruction by the exception handler is exited that will operate properly with ContinuationState≠0.

The ExceptionState bits are ephemeral because they are explicitly set by event handling and cleared by the termination of event handling, including event handling that results in a context switch.

Event Register

Events are single-bit messages used to communicate the occurrence of events between threads and interface devices.




The Event Register appears at several locations in memory, with slightly different side effects on read and write operations.

offset	side effect on read	side effect on write
0	none: return event register contents	normal: write data into event register
256	stall thread until contents of event register is non-zero, then return event register contents	stall thread until bitwise-and of data and event register contents is non-zero
512	return zero value (so read-modify-write for byte/doublet/quadlet store works)	one bits in data set (to one) corresponding event register bits
768	return zero value (so read-modify-write for byte/doublet/quadlet store works)	one bits in data clear (to zero) corresponding event register bits

Physical address

The Event Register appears at three different locations, for which three functions of the Event Register are performed as described above. The physical address of an Event Register for function f, byte b is:



Definition

```

def data ← AccessPhysicalEventRegister(pa,op,wdata) as
    f ← pa9..8
    if (pa23..10 = 0) and (pa7..4 = 0) and (f ≠ 1) then
        case f II op of
            0 II R:
                data ← 064 II EventRegister
            2 II R, 3 II R:
                data ← 0
            0 II W:
                EventRegister ← wdata63..0
            2 II W:
                EventRegister ← EventRegister or wdata63..0
            3 II W:
                EventRegister ← EventRegister and ~wdata63..0
        endcase
    else
        data ← 0
    endif
enddef

```

Events:

The table below shows the events and their corresponding event number. The priority of these events is soft, in that dispatching from the event register is controlled by software.

TODO notwithstanding the above, using the E.LOGMOST.U instruction is handy for prioritizing these events, so if you've got a preference as to numbering, speak up!

number	event
0	Clock
1	A20M# active
2	BF0 active
3	BF1 active
4	BF2 active
5	BUSCHK# active
6	FLUSH# active
7	FRCMC# active
8	IGNNE# active
9	INIT active
10	INTR active
11	NMI active
12	SMI# active
13	STPCLK# active
14	CPUTYP active at reset (Primary vs Dual processor)
15	DPEN#active at reset (Dual processing enable – driven low by dual processor)
16	FLUSH# active at reset (tristate test mode)
17	INIT active at reset
18	Bus lock broken
19	BRYRC# active at reset (drive strength)
20	

Event Mask

The Event Mask (one per thread) control whether each of the events described above is permitted to cause an exception in the corresponding thread.

Physical address

There are as many Event Masks as threads. The physical address of an Event Mask for thread th, byte b is:

63		2423 1918		32 0
FFFF FFFF 0E00 0000	63..24	th	0	b
40		5	16	3

Definition

```

def data ← AccessPhysicalEventMask(pa,op,wdata) as
    th ← pa23..19
    if {th < T} and {pa18..4 = 0} then
        case op of
            R:
                data ← 064 || EventMask[th]
            W:
                EventMask[th] ← wdata63..0
        endcase
    else
        data ← 0
    endif
enddef

```

Exceptions:

The table below shows the exceptions, the corresponding exception number, and the parameter supplied by the exception handler in register 3.

number	exception	parameter (register 3)
0	EventInterrupt	
1	MissInGlobalTB	global address
2	AccessDetailRequiredByTag	global address
3	AccessDetailRequiredByGlobalTB	global address
4	AccessDetailRequiredByLocalTB	local address
5		
6	SecondException	
7	ReservedInstruction	instruction
8	AccessDisallowedByVirtualAddress	local address
9	AccessDisallowedByTag	global address
10	AccessDisallowedByGlobalTB	global address
11	AccessDisallowedByLocalTB	local address
12	MissInLocalTB	local address
13	FixedPointArithmetic	instruction
14	FloatingPointArithmetic	instruction
15	GatewayDisallowed	none
16		
17		
18		
19		
20		
21		
22		
23		
24		
25		
	<i>TakenBranch</i>	
	<i>TakenBranchContinue</i>	

GlobalTBMiss Handler

The GlobalTBMiss exception occurs when a load, store, or instruction fetch is attempted while none of the GlobalTB entries contain a matching virtual address. The Zeus processor uses a fast software-based exception handler to fill in a missing GlobalTB entry.

There are several possible ways that software may maintain page tables. For purposes of this discussion, it is assumed that a virtual page table is maintained, in which 128 bit GTB values for each 4k byte page in a linear table which is itself in virtual memory. By maintaining the page table in virtual memory, very large virtual spaces may be managed without keeping a large amount of physical memory dedicated to page tables.

Because the page table is kept in virtual memory, it is possible that a valid reference may cause a second GTBMiss exception if the virtual address that contains the page table is not present in the GTB. The processor is designed to permit a second exception to occur within an exception handler, causing a branch to the SecondException handler. However, to simplify the hardware involved, a SecondException exception saves no specific information about the exception – handling depends on keeping enough relevant information in registers to recover from the second exception.

Zeus is a multithreaded processor, which creates some special considerations in the exception handler. Unlike a single-threaded processor, it is possible that multiple threads may nearly simultaneously reference the same page and invoke two or more GTB misses, and the fully-associative construction of the GTB requires that there be no more than one matching entry for each global virtual address. Zeus provides a search-and-insert operation (GTBUpdateFill) to simplify the handling of the GTB. This operation also uses hardware GTB pointer registers to select GTB entries for replacement in FIFO priority.

A further problem is that software may need to modify the protection information contained in the GTB, such as to remove read and/or write access to a page in order to infer which parts of memory are in use, or to remove pages from a task. These modifications may occur concurrently with the GTBMiss handler, so software must take care to properly synchronize these operations. Zeus provides a search-and-update operation (GTBUpdate) to simplify updating GTB entries.

When a large number of page table entries must be changed, noting the limited capacity of the GTB can reduce the work. Reading the GTB can be less work than matching all modified entries against the GTB contents. To facilitate this, Zeus also provides read access to the hardware GTB pointers to further permit scanning the GTB for entries which have been replaced since a previous scan. GTB pointer wraparound is also logged, so it can be determined that the entire GTB needs to be scanned if all entries have been replaced since a previous scan.

In the code below, offsets from r1 are used with the following data structure

Offset	Meaning
0..15	r0 save
16..32	r1 save
32..47	r2 save
48..63	r3 save
512..527	r4 save
528..535	BasePT
536..543	GTBUpdateFill
544..559	DummyPT
560..639	available 96 bytes

$$\text{BasePT} = 512 + 16$$

$$\text{GTBUpdateFill} = \text{BasePT} + 8$$

$$\text{DummyPT} = \text{GTBUpdateFill} + 8$$

On a GTBMiss, the handler retrieves a base address for the virtual page table and constructs an index by shifting away the page offset bits of the virtual address. A single 128-bit indexed load retrieves the new GTB entry directly (except that a virtual page table miss causes a second exception, handled below). A single 128-bit store to the GTBUpdateFill location places the entry into the GTB, after checking to ensure that a concurrent handler has not already placed the entry into the GTB.

Code for GlobalTBMiss:

```

    li64la      r2=r1,BasePT      //base address for page table
    ashri      r3@12            //4k pages
    l128la      r3=r2,r3        //retrieve page table, SecExc if bad va
2:   li64la      r2=r1,GTBUpdateFill //pointer to GTB update location
    si128la    r3,r2,0          //save new TB entry
    li128la    r3=r1,48         //restore r3
    li128la    r2=r1,32         //restore r2
    li128la    r1=r1,16         //restore r1
    bback                  //restore r0 and return

```

A second exception occurs on a virtual page table miss. It is possible to service such a page table miss directly, however, the page offset bits of the virtual address have been shifted away, and have been lost. These bits can be recovered: in such a case, a dummy GTB entry is constructed, which will cause an exception other than GTBMiss upon returning. A re-execution of the offending code will then invoke a more extensive handler, making the full virtual address available.

For purposes of this example, it is assumed that checking the contents of r2 against the contents of BasePT is a good way to ensure that the second exception handler was entered from the GlobalTBMiss handler.

Code for SecondException:

```

    si128la    r4,r1,512        //save r4
    li64la      r4=r1,BasePT      //base address for page table
    bne       r2,r4,1f          //did we lose at page table load?

```

```

li128la      r2=r1,DummyPT    //dummy page table, shifted left 64-12 bits
xshlmi128   r3@r2,64+12     //combine page number with dummy entry
li128la      r4=r1,512      //restore r4
b             2b              //fall back into GTB Miss handler
1:

```

Exceptions in detail

There are no special registers to indicate details about the exception, such as the virtual address at which an access was attempted, or the operands of a floating-point operation that results in an exception. Instead, this information is available via general-purpose registers or registers stored in memory.

When a synchronous exception or asynchronous event occurs, the original contents of registers 0..3 are saved in memory and replaced with (0) program counter, privilege level, and ephemeral program state, (1) event data pointer, (2) exception code, and (3) when applicable, failing address or instruction. A new program counter and privilege level is loaded from memory and execution begins at the new address. After handling the exception and restoring all but one register, a branch-back instruction restores the final register and resumes execution.

During exception handling, any asynchronous events are kept pending until a BranchBack instruction is performed. By this mechanism, we can handle exceptions and events one at a time, without the need to interrupt and stack exceptions. Software should take care to avoid keeping the handling of asynchronous events pending for too long.

When a second exception occurs in a thread which is handling an exception, all the above operations occur, except for the saving and replacing of registers 0..3 in memory. A distinct exception code SecondException replaces the normal exception code. By this mechanism, a fast exception handler for GlobalTBMiss can be written, in which a second GlobalTBMiss or FixedPointOverflow exception may safely occur.

When a third exception occurs in a thread which is handling an exception, an immediate transfer of control occurs to the machine check vector address, with information about the exception available in the **machine check cause** field of the **status register**. The transfer of control may overwrite state that may be necessary to recover from the exception; the intent is to provide a satisfactory post-mortem indication of the characteristics of the failure.

This section describes in detail the conditions under which exceptions occur, the parameters passed to the exception handler, and the handling of the result of the procedure.

Reserved Instruction

The ReservedInstruction exception occurs when an instruction code which is reserved for future definition as part of the Zeus architecture is executed.

Register 3 contains the 32-bit instruction.

Access Disallowed by virtual address

This exception occurs when a load, store, branch, or gateway refers to an aligned memory operand with an improperly aligned address, or if architecture description parameter **LB=1**, may also occur if the add or increment of the base register or program counter which generates the address changes the unmasked upper 16 bits of the local address.

Register 3 contains the local address to which the access was attempted.

Access disallowed by tag

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which the matching cache tag entry does not permit this access.

Register 3 contains the global address to which the access was attempted.

Access detail required by tag

This exception occurs when a read (load), write (store), or execute attempts to access a virtual address for which the matching virtual cache entry would permit this access, but the detail bit is set.

Register 3 contains the global address to which the access was attempted.

Description

The exception handler should determine accessibility. If the access should be allowed, the continuepastdetail bit is set and execution returns. Upon return, execution is restarted and the access will be retried. Even if the detail bit is set in the matching virtual cache entry, access will be permitted.

Access disallowed by global TB

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which the matching global TB entry does not permit this access.

Register 3 contains the global address to which the access was attempted.

Description

The exception handler should determine accessibility, modify the virtual memory state if desired, and return if the access should be allowed. Upon return, execution is restarted and the access will be retried.

Access detail required by global TB

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which the matching global TB entry would permit this access, but the detail bit in the global TB entry is set.

Register 3 contains the global address to which the access was attempted.

Description

The exception handler should determine accessibility and return if the access should be allowed. Upon return, execution is restarted and the access will be allowed. If the access is not to be allowed, the handler should not return.

Global TB miss

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which no global TB entry matches.

Register 3 contains the global address to which the access was attempted.

Description

The exception handler should load a global TB entry that defines the translation and protection for this address. Upon return, execution is restarted and the global TB access will be attempted again.

Access disallowed by local TB

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which the matching local TB entry does not permit this access.

Register 3 contains the local address to which the access was attempted.

Description

The exception handler should determine accessibility, modify the virtual memory state if desired, and return if the access should be allowed. Upon return, execution is restarted and the access will be retried.

Access detail required by local TB

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which the matching local TB entry would permit this access, but the detail bit in the local TB entry is set.

Register 3 contains the local address to which the access was attempted.

Description

The exception handler should determine accessibility and return if the access should be allowed. Upon return, execution is restarted and the access will be allowed. If the access is not to be allowed, the handler should not return.

Local TB miss

This exception occurs when a read (load), write (store), execute, or gateway attempts to access a virtual address for which no local TB entry matches.

Register 3 contains the local address to which the access was attempted.

Description

The exception handler should load a local TB entry that defines the translation and protection for this address. Upon return, execution is restarted and the local TB access will be attempted again.

Floating-point arithmetic

Register 3 contains the 32-bit instruction.

Description

The address of the instruction that was the cause of the exception is passed as the contents of register 0. The exception handler should attempt to perform the function specified in the instruction and service any exceptional conditions that occur.

Fixed-point arithmetic

Register 3 contains the 32-bit instruction.

Description

The address of the instruction which was the cause of the exception is passed as the contents of register 0. The exception handler should attempt to perform the function specified in the instruction and service any exceptional conditions that occur.

Reset and Error Recovery

Certain external and internal events cause the processor to invoke reset or error recovery operations. These operations consist of a full or partial reset of critical machine state, including initialization of the threads to begin fetching instructions from the start vector address. Software may determine the nature of the reset or error by reading the value of the control register, in which finding the **reset** bit set (1) indicates that a reset has occurred, and finding both the **reset** bit cleared (0) indicates that a machine check has occurred. When either a reset or machine check has been indicated, the contents of the status register contain more detailed information on the cause.

Definition

```
def PerformMachineCheck(cause) as
    ResetVirtualMemory()
    ProgramCounter ← StartVectorAddress
    PrivilegeLevel ← 3
    StatusRegister ← cause
enddef
```

Reset

A reset may be caused by a power-on reset, a bus reset, a write of the control register which sets the **reset** bit, or internally detected errors including meltdown detection, and double check.

A reset causes the processor to set the configuration to minimum power and low clock speed, note the cause of the reset in the status register, stabilize the phase locked loops, disable the MMU from the control register, and initialize all threads to begin execution at the start vector address.

Other system state is left undefined by reset and must be explicitly initialized by software; this explicitly includes the thread register state, LTB and GTB state, superspring state, and external interface devices. The code at the start vector address is responsible for initializing these remaining system facilities, and reading further bootstrap code from an external ROM.

Power-on Reset

A reset occurs upon initial power-on. The cause of the reset is noted by initializing the Status Register and other registers to the reset values noted below.

Bus Reset

A reset occurs upon observing that the RESET signal has been at active. The cause of the reset is noted by initializing the Status Register and other registers to the reset values noted below.

Control Register Reset

A reset occurs upon writing a one to the **reset** bit of the Control Register. The cause of the reset is noted by initializing the Status Register and other registers to the reset values noted below.

Meltdown Detected Reset

A reset occurs if the temperature is above the threshold set by the **meltdown margin** field of the configuration register. The cause of the reset is noted by setting the **meltdown detected** bit of the Status Register.

Double Check Reset

A reset occurs if a second machine check occurs that prevents recovery from the first machine check. Specifically, the occurrence of an **exception in event thread, watchdog timer error, or bus error** while any machine check cause bit is still set in the Status Register results in a double machine check reset. The cause of the reset is noted by setting the **double check** bit of the Status Register.

Machine Check

Detected hardware errors, such as communications errors in the bus, a watchdog timeout error, or internal cache parity errors, invoke a machine check. A machine check will disable the MMU, to translate all local virtual addresses to equal physical addresses, note the cause of the exception in the Status Register, and transfer control of the all threads to the start vector address. This action is similar to that of a reset, but differs in that the configuration settings, and thread state are preserved.

Recovery from machine checks depends on the severity of the error and the potential loss of information as a direct cause of the error. The start vector address is designed to reach internal ROM memory, so that operation of machine check diagnostic and recovery code need not depend on proper operation or contents of any external device. The program counter and register file state of the thread prior to the machine check is lost (except for the portion of the program counter saved in the Status Register), so diagnostic and recovery code must not assume that the register file state is indicative of the prior operating state of the thread. The state of the thread is frozen similarly to that of an exception.

Machine check diagnostic code determines the cause of the machine check from the processor's Status Register, and as required, the status and other registers of external bus devices.

Recovery code will generally consume enough time that real-time interface performance targets may have been missed. Consequently, the machine check recovery software may need to repair further damage, such as interface buffer underruns and overruns as may have occurred during the intervening time.

This final recovery code, which re-initializes the state of the interface system and recovers a functional event thread state, may return to using the complete machine resources, as the condition which caused the machine check will have been resolved.

The following table lists the causes of machine check errors.

Parity or uncorrectable error in on-chip cache
Parity or communications error in system bus
Event Thread exception
Watchdog timer

machine check errors

Parity or Uncorrectable Error in Cache

When a parity or uncorrectable error occurs in an on-chip cache, such an error is generally non-recoverable. These errors are non-recoverable because the data in such caches may reside anywhere in memory, and because the data in such caches may be the only up-to-date copy of that memory contents. Consequently, the entire contents of the memory store is lost, and the severity of the error is high enough to consider such a condition to be a system failure.

The machine check provides an opportunity to report such an error before shutting down a system for repairs.

There are specific means by which a system may recover from such an error without failure, such as by restarting from a system-level checkpoint, from which a consistent memory state can be recovered.

Parity or Communications Error in Bus

When a parity or communications error occurs in the system bus, such an error may be partially recoverable.

Bits corresponding to the affected bus operation are set in the processor's Status Register. Recovery software should determine which devices are affected, by querying the Status Register of each device on the affected MediaChannel channels.

A bus timeout may result from normal self-configuration activities.

If the error is simply a communications error, resetting appropriate devices and restarting tasks may recover from the error. Read and write transactions may have been underway at the time of a machine check and may or may not be reflected in the current system state.

If the error is from a parity error in memory, the contents of the affected area of memory is lost, and consequently the tasks associated with that memory must generally be aborted, or resumed from a task-level checkpoint. If the contents of the affected memory can be recovered from mass storage, a complete recovery is possible.

If the affected memory is that of a critical part of the operating system, such a condition is considered a system failure, unless recovery can be accomplished from a system-level checkpoint.

Watchdog Timeout Error

A watchdog timeout error indicates a general software or hardware failure. Such an error is generally treated as non-recoverable and fatal.

Event Thread Exception

When an event thread suffers an exception, the cause of the exception and a portion of the virtual address at which the exception occurred are noted in the Status Register. Because under normal circumstances, the event thread should be designed not to encounter exceptions, such exceptions are treated as non-recoverable, fatal errors.

Reset state

A reset or machine check causes the Zeus processor to stabilize the phase locked loops, disable the local and global TB, to translate all local virtual addresses to equal physical addresses, and initialize all threads to begin execution at the start vector address.

Start Address

The start address is used to initialize the threads with a program counter upon a reset, or machine check. These causes of such initialization can be differentiated by the contents of the Status Register.

The start address is a virtual address which, when “translated” by the local TB and global TB to a physical address, is designed to access the internal ROM code. The internal ROM space is chosen to minimize the number of internal resources and interfaces that must be operated to begin execution or recover from a machine check.

Virtual/physical address	description
0xFFFF FFFF FFFF FFFC	start vector address

Definition

```
def StartProcessor as
    forever
        catch check
            EnableWatchdog ← 0
            fork RunClock
            ControlRegister62 ← 0
            for th ← 0 to T-1
                ProgramCounter[th] ← 0xFFFF FFFF FFFF FFFC
                PrivilegeLevel[th] ← 3
                fork Thread(th)
            endfor
        endcatch
    endforever
```

```
endcatch
kill RunClock
for th ← 0 to T-1
    kill Thread(th)
endfor
PerformMachineCheck(check)
endforever
enddef

def PerformMachineCheck(check) as
    case check of
        ClockWatchdog:
        CacheError:
        ThirdException:
    endcase
enddef
```

Internal ROM Code

Zeus internal ROM code performs reset initialization of on-chip resources, including the LZC and LOC, followed by self-testing. The BIOS ROM should be scanned for a special prefix that indicates that Zeus native code is present in the ROM, in which case the ROM code is executed directly, otherwise execution of a BIOS-level x86 emulator is begun.

Memory and Devices

Physical Memory Map

Zeus defines a 64-bit physical address, but while residing in a S7 pin-out, can address a maximum of 4Gb of main memory. In other packages the core Zeus design can provide up to 64-bit external physical address spaces. Bit 63..32 of the physical address distinguishes between internal (on-chip) physical addresses, where bits 63..32=FFFFFF, and external (off-chip) physical addresses, where bits 63..32 \neq FFFFFF.

Address range	bytes	Meaning
0000 0000 0000 0000..0000 0000 FFFF FFFF	4G	External Memory
0000 0001 0000 0000..FFFF FFFE FFFF FFFF	16E-8G	External Memory expansion
FFFF FFFF 0000 0000..FFFF FFFF 0002 0FFF	128K+4K	Level One Cache
FFFF FFFF 0002 1000..FFFF FFFF 08FF FFFF	144M-132K	Level One Cache expansion
FFFF FFFF 0900 0000..FFFF FFFF 0900 007F	128	Level One Cache redundancy
FFFF FFFF 0900 0080..FFFF FFFF 09FF FFFF	16M-128	LOC redundancy expansion
FFFF FFFF 0A00 0000+t*2 ^{19+e*16}	8*T*2 ^{LE}	LTB thread t entry e
FFFF FFFF 0A00 0000..FFFF FFFF 0AFF FFFF	8*T*2 ^{LE}	LTB max 8*T*2 ^{LE} = 16M bytes
FFFF FFFF 0B00 0000..FFFF FFFF 0BFF FFFF	16M	Special Bus Operations
FFFF FFFF 0C00 0000+t _{5..GT} *2 ^{19+GT+e*16}	T2 ^{4+GE-GT}	GTB thread t entry e
FFFF FFFF 0C00 0000..FFFF FFFF 0CFF FFFF	T2 ^{4+GE-GT}	GTB max 2 ⁵⁺⁴⁺¹⁵ = 16M bytes
FFFF FFFF 0D00 0000+t _{5..GT} *2 ^{19+GT}	16*T*2 ^{GT}	GTBUpdate thread t
FFFF FFFF 0D00 0100+t _{5..GT} *2 ^{19+GT}	16*T*2 ^{GT}	GTBUpdateFill thread t
FFFF FFFF 0D00 0200+t _{5..GT} *2 ^{19+GT}	8*T*2 ^{GT}	GTBLast thread t
FFFF FFFF 0D00 0300+t _{5..GT} *2 ^{19+GT}	8*T*2 ^{GT}	GTBFirst thread t
FFFF FFFF 0D00 0400+t _{5..GT} *2 ^{19+GT}	8*T*2 ^{GT}	GTBBump thread t
FFFF FFFF 0E00 0000+t*2 ¹⁹	8T	Event Mask thread t
FFFF FFFF 0F00 0008..FFFF FFFF 0F00 00FF	256-8	Reserved
FFFF FFFF 0F00 0100..FFFF FFFF 0F00 0107	8	Event Register with stall
FFFF FFFF 0F00 0108..FFFF FFFF 0F00 01FF	256-8	Reserved
FFFF FFFF 0F00 0200..FFFF FFFF 0F00 0207	8	Event Register bit set
FFFF FFFF 0F00 0208..FFFF FFFF 0F00 02FF	256-8	Reserved
FFFF FFFF 0F00 0300..FFFF FFFF 0F00 0307	8	Event Register bit clear
FFFF FFFF 0F00 0308..FFFF FFFF 0F00 03FF	256-8	Reserved
FFFF FFFF 0F00 0400..FFFF FFFF 0F00 0407	8	Clock Cycle
FFFF FFFF 0F00 0408..FFFF FFFF 0F00 04FF	256-8	Reserved
FFFF FFFF 0F00 0500..FFFF FFFF 0F00 0507	8	Thread
FFFF FFFF 0F00 0508..FFFF FFFF 0F00 05FF	256-8	Reserved
FFFF FFFF 0F00 0600..FFFF FFFF 0F00 0607	8	Clock Event
FFFF FFFF 0F00 0608..FFFF FFFF 0F00 06FF	256-8	Reserved
FFFF FFFF 0F00 0700..FFFF FFFF 0F00 0707	8	Clock Watchdog
FFFF FFFF 0F00 0708..FFFF FFFF 0F00 07FF	256-8	Reserved
FFFF FFFF 0F00 0800..FFFF FFFF 0F00 0807	8	Tally Counter 0
FFFF FFFF 0F00 0808..FFFF FFFF 0F00 08FF	256-8	Reserved
FFFF FFFF 0F00 0900..FFFF FFFF 0F00 0907	8	Tally Control 0
FFFF FFFF 0F00 0908..FFFF FFFF 0F00 09FF	256-8	Reserved
FFFF FFFF 0F00 0A00..FFFF FFFF 0F00 0A07	8	Tally Counter 1

FFFF FFFF OF00 0A08..FFFF FFFF OF00 0AFF	256-8	Reserved
FFFF FFFF OF00 0B00..FFFF FFFF OF00 0B07	8	Tally Control 1
FFFF FFFF OF00 0B08..FFFF FFFF OF00 0BFF	256-8	Reserved
FFFF FFFF OF00 0C00..FFFF FFFF OF00 0C07	8	Exception Base
FFFF FFFF OF00 0C08..FFFF FFFF OF00 0CFF	512-8	Reserved
FFFF FFFF OF00 0C00..FFFF FFFF OF00 0D07	8	Bus Control Register
FFFF FFFF OF00 0D08..FFFF FFFF OF00 0DFF	512-8	Reserved
FFFF FFFF OF00 0E00..FFFF FFFF OF00 0E07	8	Status Register
FFFF FFFF OF00 0208..FFFF FFFF OF00 02FF	256-8	Reserved
FFFF FFFF OF00 0F00..FFFF FFFF OF00 0F07	8	Control Register
FFFF FFFF OF00 0F08..FFFF FFFF FEFF FFFF		Reserved
FFFF FFFF FF00 0000..FFFF FFFF FFFE FFFF	16M-64k	Internal ROM expansion
FFFF FFFF FFFF 0000..FFFF FFFF FFFF FFFF	64K	Internal ROM

The suffixes in the table above have the following meanings:

letter	name	2 ^x	"binary"	10 ^y	"decimal"
b	bits				
B	bytes	0	1	0	1
K	kilo	10	1 024	3	1 000
M	mega	20	1 048 576	6	1 000 000
G	giga	30	1 073 741 824	9	1 000 000 000
T	tera	40	1 099 511 627 776	12	1 000 000 000 000
P	peta	50	1 125 899 906 842 624	15	1 000 000 000 000 000
E	exa	60	1 152 921 504 606 846 976	18	1 000 000 000 000 000 000

Definition

```

def data ← ReadPhysical(pa,size) as
    data,flags ← AccessPhysical(pa,size,W,A,R,0)
enddef

def WritePhysical(pa,size,wdata) as
    data,flags ← AccessPhysical(pa,size,W,A,W,wdata)
enddef

def data,flags ← AccessPhysical(pa,size,cc,op,wdata) as
    if {0x0000000000000000 ≤ pa ≤ 0x00000000FFFFFF} then
        data,flags ← AccessPhysicalBus(pa,size,cc,op,wdata)
    else
        data ← AccessPhysicalDevices(pa,size,op,wdata)
        flags ← 1
    endif
enddef

def data ← AccessPhysicalDevices(pa,size,op,wdata) as
    if {size=256} then
        data0 ← AccessPhysicalDevices(pa,128.op.wdata127..0)
        data1 ← AccessPhysicalDevices(pa+16,128.op.wdata255..128)
        data ← data1 || data0
    elseif {0xFFFFFFFF0B000000 ≤ pa ≤ 0xFFFFFFFF0BFFFFFF} then
        //don't perform RMW on this region
        data ← AccessPhysicalOtherBus(pa,size,op,wdata)
    endif
enddef

```

```

elseif {op=W} and {size<128} then
    //this code should change to check pa4..0≠0 and size<sizeofreg
    rdata ← AccessPhysicalDevices(pa and -15,128,R,0)
    bs ← 8*(pa and 15)
    be ← bs + size
    hdata ← rdata[127..be] || wdata[be-1..bs] || rdata[bs-1..0]
    data ← AccessPhysicalDevices(pa and -15,128,W,hdata)
elseif {0x0000000010000000 ≤ pa ≤ 0xFFFFFFFFFFFFFFFFFF} then
    data ← 0
elseif {0xFFFFFFFF00000000 ≤ pa ≤ 0xFFFFFFFF08FFFFFF} then
    data,← AccessPhysicalLOC(pa,op,wdata)
elseif {0xFFFFFFFF09000000 ≤ pa ≤ 0xFFFFFFFF09FFFFFF} then
    data ← AccessPhysicalLOCRedundancy(pa,op,wdata)
elseif {0xFFFFFFFF0A000000 ≤ pa ≤ 0xFFFFFFFF0AFFFFFF} then
    data ← AccessPhysicalLTB(pa,op,wdata)
elseif {0xFFFFFFFF0C000000 ≤ pa ≤ 0xFFFFFFFF0CFFFFFF} then
    data ← AccessPhysicalGTB(pa,op,wdata)
elseif {0xFFFFFFFF0D000000 ≤ pa ≤ 0xFFFFFFFF0DFFFFFF} then
    data ← AccessPhysicalGTBRegisters(pa,op,wdata)
elseif {0xFFFFFFFF0E000000 ≤ pa ≤ 0xFFFFFFFF0EFFFFFF} then
    data ← AccessPhysicalEventMask(pa,op,wdata)
elseif {0xFFFFFFFF0F000000 ≤ pa ≤ 0xFFFFFFFF0FFFFFFF} then
    data ← AccessPhysicalSpecialRegisters(pa,op,wdata)
elseif {0xFFFFFFFF10000000 ≤ pa ≤ 0xFFFFFFFFEFFFFFFF} then
    data ← 0
elseif {0xFFFFFFFFFFF000000 ≤ pa ≤ 0xFFFFFFFFFFFFFFFFF} then
    data ← AccessPhysicalROM(pa,op,wdata)
endif
enddef

def data ← AccessPhysicalSpecialRegisters(pa,op,wdata) as
    if {pa7..0 ≥ 0x10} then
        data ← 0
    elseif {0xFFFFFFFF0F000000 ≤ pa ≤ 0xFFFFFFFF0F0003FF} then
        data ← AccessPhysicalEventRegister(pa,op,wdata)
    elseif {0xFFFFFFFF0F000500 ≤ pa ≤ 0xFFFFFFFF0F0005FF} then
        data,← AccessPhysicalThread(pa,op,wdata)
    elseif {0xFFFFFFFF0F000400 ≤ pa ≤ 0xFFFFFFFF0F0007FF} then
        data,← AccessPhysicalClock(pa,op,wdata)
    elseif {0xFFFFFFFF0F000800 ≤ pa ≤ 0xFFFFFFFF0F000BFF} then
        data,← AccessPhysicalTally(pa,op,wdata)
    elseif {0xFFFFFFFF0F000C00 ≤ pa ≤ 0xFFFFFFFF0F000CFF} then
        data,← AccessPhysicalExceptionBase(pa,op,wdata)
    elseif {0xFFFFFFFF0F000D00 ≤ pa ≤ 0xFFFFFFFF0F000DFF} then
        data,← AccessPhysicalBusControl(pa,op,wdata)
    elseif {0xFFFFFFFF0F000E00 ≤ pa ≤ 0xFFFFFFFF0F000EFF} then
        data,← AccessPhysicalStatus(pa,op,wdata)
    elseif {0xFFFFFFFF0F000F00 ≤ pa ≤ 0xFFFFFFFF0F000FFF} then
        data,← AccessPhysicalControl(pa,op,wdata)
    endif
enddef

```

Architecture Description Register

The last hexlet of the internal ROM contains data that describes implementation-dependent choices within the architecture specification. The last quadlet of the internal ROM contains a branch-immediate instruction, so the architecture description is limited to 96 bits.

Address range	bytes	Meaning
FFFF FFFF FFFF FFFC..FFFF FFFF FFFF FFFF	4	Reset address
FFFF FFFF FFFF FFF0..FFFF FFFF FFFF FFFB	12	Architecture Description Register

The table below indicates the detailed layout of the Architecture Description Register.

bits	field name	value range	interpretation
127..96	bl start		Contains a branch instruction for bootstrap from internal ROM
95..23	0	0	reserved
22..21	GT	1	log ₂ threads which share a global TB
20..17	GE	7	log ₂ entries in global TB
16	LB	1	local TB based on base register
15..14	LE	1	log ₂ entries in local TB (per thread)
13	CT	1	dedicated tags in first-level cache
12..10	CS	2	log ₂ cache blocks in first-level cache set
9..5	CE	9	log ₂ cache blocks in first-level cache
4..0	T	4	number of execution threads

The architecture description register contains a machine-readable version of the architecture framework parameters: **T**, **CE**, **CS**, **CT**, **LE**, **GE**, and **GT** described in the [Architectural Framework](#) section on page 17.

Status Register

The status register is a 64-bit register with both read and write access, though the only legal value which may be written is a zero, to clear the register. The result of writing a non-zero value is not specified.

bits	field name	value range		interpretation
63	power-on	1	0..1	This bit is set when a power-on reset has caused a reset.
62	internal r set	0	0..1	This bit is set when writing to the control register caused a reset.
61	bus reset	0	0..1	This bit is set when a bus reset has caused a reset.
60	double check	0	0..1	This bit is set when a double machine check has caused a reset.
59	meltdown	0	0..1	This bit is set when the meltdown detector has caused a reset.
58..56	0	0*	0	Reserved for other machine check causes.
55	event exception	0	0..1	This bit is set when an exception in event thread has caused a machine check.
54	watchdog timeout	0	0..1	This bit is set when a watchdog timeout has caused a machine check.
53	bus error	0	0..1	This bit is set when a bus error has caused a machine check.
52	cache error	0	0..1	This bit is set when a cache error has caused a machine check.
51	vm error	0	0..1	This bit is set when a virtual memory error has caused a machine check.
50..48	0	0*	0	Reserved for other machine check causes.
47..32	machine check detail	0*	0..40 95	Set to exception code if Exception in event thread. Set to bus error code if bus error.
31..0	machine check program counter	0	0	Set to indicate bits 31..0 of the value of the thread 0 program counter at the initiation of a machine check.

The **power-on** bit of the status register is set upon the completion of a power-on reset.

The **bus reset** bit of the status register is set upon the completion of a bus reset initiated by the RESET pin of the Socket 7 interface.

The **double check** bit of the status register is set when a second machine check occurs that prevents recovery from the first machine check, or which is indicative of machine check recovery software failure. Specifically, the occurrence of an **event exception**, **watchdog timeout**, **bus error**, or **meltdown** while any reset or machine check cause bit of the status register is still set results in a **double check** reset.

The **meltdown** bit of the status register is set when the meltdown detector has discovered an on-chip temperature above the threshold set by the **meltdown threshold** field of the control register, which causes a reset to occur.

The **event exception** bit of the status register is set when an event thread suffers an exception, which causes a machine check. The exception code is loaded into the **machine**

check detail field of the status register, and the machine check program counter is loaded with the low-order 32 bits of the program counter and privilege level.

The **watchdog timeout** bit of the status register is set when the **watchdog timer** register is equal to the **clock cycle** register, causing a machine check.

The **bus error** bit of the status register is set when a bus transaction error (bus timeout, invalid transaction code, invalid address, parity errors) has caused a machine check.

The **cache error** bit of the status register is set when a cache error, such as a cache parity error has caused a machine check.


The **vm error** bit of the status register is set when a virtual memory error, such as a GTB multiple-entry selection error has caused a machine check.

The **machine check detail** field of the status register is set when a machine check has been completed. For an exception in event thread, the value indicates the type of exception for which the most recent machine check has been reported. For a bus error, this field may indicate additional detail on the cause of the bus error. For a cache error, this field may indicate the address of the error at which the cache parity error was detected

The **machine check program counter** field of the status register is loaded with bits 31..0 of the program counter and privilege level at which the most recent machine check has occurred. The value in this field provides a limited diagnostic capability for purposes of software development, or possibly for error recovery.

Physical address

The physical address of the Status Register, byte b is:



Definition

```
def data ← AccessPhysicalStatus(pa,op,wdata) as
  case op of
    R:
      data ← 064 || StatusRegister
    W:
      StatusRegister ← wdata63..0
  endcase
enddef
```

Control Register

The control register is a 64-bit register with both read and write access. It is altered only by write access to this register.

bits	field name	value	range	interpretation
63	reset	0	0..1	set to invoke internal reset
62	MMU	0	0..1	set to enable the MMU
61	LOC parity	0	0..1	set to enable LOC parity
60	meltdown	0	0..1	set to enable meltdown detector
59..57	LOC timing	0	0..7	adjust LOC timing 00slow..70fast
56..55	LOC stress	0	0..3	adjust LOC stress 00normal
54..52	clock timing	0	0..7	adjust clock timing 00slow..70fast
51..12	0	0	0	Reserved
11..8	global access	0*	0..15	global access
7..0	niche limit	0*	0..12 7	niche limit

The **reset** bit of the control register provides the ability to reset an individual Zeus device in a system. Writing a one (1) to this bit is equivalent to a power-on reset or a bus reset. The duration of the reset is sufficient for the operating state changes to have taken effect. At the completion of the reset operation, the **internal reset** bit of the status register is set and the **reset** bit of the control register is cleared (0).

The **MMU** bit of the control register provides the ability to enable or disable the MMU features of the Zeus processor. Writing a zero (0) to this bit disables the MMU, causing all MMU-related exceptions to be disabled and causing all load, store, program and gateway virtual addresses to be treated as physical addresses. Writing a one (1) to this bit enables the MMU and MMU-related exceptions. On a reset or machine check, this bit is cleared (0), thus disabling the MMU.

The **parity** bit of the control register provides the ability to enable or disable the cache parity feature of the Zeus processor. Writing a zero (0) to this bit disables the parity check, causing the parity check machine check to be disabled. Writing a one (1) to this bit enables the cache parity machine check. On a reset or machine check, this bit is cleared (0), thus disabling the cache parity check.

The **meltdown** bit of the control register provides the ability to enable or disable the meltdown detection feature of the Zeus processor. Writing a zero (0) to this bit disables the meltdown detector, causing the meltdown detected machine check to be disabled. Writing a one (1) to this bit enables the meltdown detector. On a reset or machine check, this bit is cleared (0), thus disabling the meltdown detector.

The **LOC timing** bits of the control register provide the ability to adjust the cache timing of the Zeus processor. Writing a zero (0) to this field sets the cache timing to its slowest state, enhancing reliability but limiting clock rate. Writing a seven (7) to this field sets the cache timing to its fastest state, limiting reliability but enhancing performance. On a reset or machine check, this field is cleared (0), thus providing operation at low clock rate. Changing this register should be performed when the cache is not actively being operated.

The **LOC stress** bits of the control register provide the ability to stress the LOC parameters by adjusting voltage levels within the LOC. Writing a zero (0) to this field sets the cache

parameters to its normal state, enhancing reliability. Writing a non-zero value (1, 2, or 3) to this field sets the cache parameters to levels at which cache reliability is slightly compromised. The stressed parameters are used to cause LOC cells with marginal performance to fail during self-test, so that redundancy can be employed to enhance reliability. On a reset or machine check, this field is cleared (0), thus providing operation at normal parameters. Changing this register should be performed when the cache is not actively being operated.

The **clock timing** bits of the control register provide the ability to adjust the clock timing of the Zeus processor. Writing a zero (0) to this field sets the clock timing to its slowest state, enhancing reliability but limiting clock rate. Writing a seven (7) to this field sets the clock timing to its fastest state, limiting reliability but enhancing performance. On a power on reset, bus reset, or machine check, this field is cleared (0), thus providing operation at low clock rate. The internal clock rate is set to $(\text{clock timing}+1)/2 \times (\text{external clock rate})$. Changing this register should be performed along with a control register reset.

The **global access** bits of the control register determine whether a local TB miss cause an exception or treatment as a global address. A single bit, selected by the privilege level active for the access from four bit configuration register field, "Global Access," (GA) determines the result. If GApL is zero (0), the failure causes an exception, if it is one (1), the failure causes the address to be used as a global address directly.

The **niche limit** bits of the control register determine which cache lines are used for cache access, and which lines are used for niche access. For addresses $\text{pa14..8} < \text{nl}$, a 7-bit address modifier register **am** is inclusive-or'd against pa14..8 to determine the cache line. The cache modifier **am** must be set to $(1^{7-\log(128-\text{nl})} \parallel 0^{\log(128-\text{nl})})$ for proper operation. The **am** value does not appear in a register and is generated from the **nl** value.

Physical address

The physical address of the Control Register, byte **b** is:



Definition

```
def data ← AccessPhysicalControl(pa,op,wdata) as
    case op of
        R:
            data ← 064 || ControlRegister
        W:
            ControlRegister ← wdata63..0
    endcase
enddef
```


Clock

The Zeus processor provides internal clock facilities using three registers, a **clock cycle** register that increments one every cycle, a **clock event** register that sets the **clock** bit in the **event register**, and a **clock watchdog** register that invokes a **clock watchdog machine check**. These registers are memory mapped.

Clock Cycle

Each Zeus processor includes a clock that maintains processor-clock-cycle accuracy. The value of the **clock cycle** register is incremented on every cycle, regardless of the number of instructions executed on that cycle. The **clock cycle** register is 64-bits long.

For testing purposes the **clock cycle** register is both readable and writable, though in normal operation it should be written only at system initialization time; there is no mechanism provided for adjusting the value in the **clock cycle** counter without the possibility of losing cycles.




Clock Event

An event is asserted when the value in the **clock cycle** register is equal to the value in the **clock event** register, which sets the **clock** bit in the **event register**.

It is required that a sufficient number of bits be implemented in the **clock event** register so that the comparison with the **clock cycle** register overflows no more frequently than once per second. 32 bits is sufficient for a 4 GHz clock. The remaining unimplemented bits must be zero whenever read, and ignored on write. Equality is checked only against bits that are implemented in both the **clock cycle** and **clock event** registers.

For testing purposes the **clock event** register is both readable and writable, though in normal operation it is normally written to.



Clock Watchdog

A Machine Check is asserted when the value in the **clock cycle** register is equal to the value in the **clock watchdog** register, which sets the **watchdog timeout** bit in the **control register**.

A Machine Check or a Reset, of any cause including a **clock watchdog**, disables the **clock watchdog machine check**. A write to the **clock watchdog** register enables the **clock watchdog machine check**.


It is required that a sufficient number of bits be implemented in the **clock watchdog** register so that the comparison with the **clock cycle** register overflows no more frequently than once per second. 32 bits is sufficient for a 4 GHz clock. The remaining unimplemented bits must be zero whenever read, and ignored on write. Equality is checked only against bits that are implemented in both the **clock cycle** and **clock watchdog** registers.

The **clock watchdog** register is both readable and writable, though in normal operation it is usually and periodically written with a sufficiently large value that the register does not equal the value in the **clock cycle** register before the next time it is written.



Physical address

The Clock registers appear at three different locations, for which three registers of the Clock are mapped. The Clock Cycle counter is register 0, the Clock Event is register 2, and ClockWatchdog is register 3. The physical address of a Clock Register **f**, byte **b** is:



Definition

```

def data ← AccessPhysicalClock(pa,op,wdata) as
    f ← pa9..8
    case f 11 op of
        0 11 R:
            data ← 064 11 ClockCycle
        0 11 W:
            ClockCycle ← wdata63..0
        2 11 R:
            data ← 096 11 ClockEvent
        2 11 W:
            ClockEvent ← wdata31..0
        3 11 R:
            data ← 096 11 ClockWatchdog
        3 11 W:
            ClockWatchdog ← wdata31..0
            EnableWatchdog ← 1
    endcase
enddef

def RunClock as
    forever
        ClockCycle ← ClockCycle + 1
        if EnableWatchdog and (ClockCycle31..0 = ClockWatchdog31..0) then
            raise ClockWatchdogMachineCheck
        elseif (ClockCycle31..0 = ClockEvent31..0) then
            EventRegister0 ← 1

```

```

        endif
        wait
    endforever
enddef

```


Tally

Tally Counter

Each processor includes two counters that can tally processor-related events or operations. The values of the **tally counter** registers are incremented on each processor clock cycle in which specified events or operations occur. The **tally counter** registers do not signal events.


It is required that a sufficient number of bits be implemented so that the tally counter registers overflow no more frequently than once per second. 32 bits is sufficient for a 4 GHz clock. The remaining unimplemented bits must be zero whenever read, and ignored on write.

For testing purposes each of the tally counter registers are both readable and writable, though in normal operation each should be written only at system initialization time; there is no mechanism provided for adjusting the value in the event counter registers without the possibility of losing counts.



Physical address

The Tally Counter registers appear at two different locations, for which the two registers are mapped. The physical address of a Tally Counter register **f**, byte **b** is:



Tally Control

The tally counter control registers each select one metric for one of the tally counters.

63	1615	0
0	tally control 0	16
48		

63	1615	0
0	tally control 1	16
48		

Each control register is loaded with a value in one of the following formats:

15	14	13 12	thread					8	7	6	5	4	3	2	1	0
0	flag							W	E	X	G	S	L	B	A	
1	2			5				1	1	1	1	1	1	1	1	

flag	meaning
0	count instructions issued
1	count instructions retired (differs by branch mispred, exceptions)
2	count cycles in which at least one instruction is issued
3	count cycles in which next instruction is waiting for issue

W E X G S L B A: include instructions of these classes

15	14	13 12	thread					87	flag					4	3	2	1	0
1	0	0						4	S	L	W	I		1	1	1	1	1
1	1	1		5														

flag	meaning
0	count bytes transferred cache/buffer to/from processor
1	count bytes transferred memory to/from cache/buffer
2	
3	
4	count cache hits
5	count cycles in which at least one cache hit occurs
6	count cache misses
7	count cycles in which at least one cache miss occurs
8..15	

S L W I: include instructions of these classes (Store, Load, Wide, Instruction fetch)

15	14	13	12		87	6	5	4	3	2	1	0
1	0	1		thread	5	flag	n	E	X	T	G	A
1	1	1				2	1	1	1	1	1	1

flag	meaning
0	count cycles in which a new instruction is issued
1	count cycles in which an execution unit is busy
2	
3	count cycles in which an instruction is waiting for issue

n select unit number for G or A unit

E X T G A: include units of these classes (Ensemble, Crossbar, Translate, Group, Address)

15	14	13	12	11	10	9	8	7	6	5		0
1	1	0	0	0	0	0	0	0	0	0		event
1	1	1	1	1	1	1	1	1	1	1		6

event: select event number from event register

15	14	13	12	11	10	9	8	7	6	5		0
1	1	1	0	0	0	0	0	0	0	0		other
1	1	1	1	1	1	1	1	1	1	1		6

Other valid values for the tally control fields are given by the following table:

other	meaning
0	count number of instructions waiting to issue each cycle
1	count number of instructions waiting in spring each cycle
2..63	Reserved

tally control field interpretation

Physical address

The Tally Control registers appear at two different locations, for which the two registers are mapped. The physical address of a Tally Control register f, byte b is:

63		1098	32	0
	FFFF FFFF 0F00 0900 63..10		f	32 b
	54		1	6 3

Definition

```
def data ← AccessPhysicalTally(pa,op,wdata) as
  f ← pa9
  case pag 11 op of
    0 11 R:
      data ← 096 11 TallyCounter[f]
    0 11 W:
      TallyCounter[f] ← wdata31..0
```

```

1 11 R:
  data ← 012 11 TallyControl[f]
1 11 W:
  TallyControl[f]← wdata15..0
endcase
enddef

```

Thread Register


The Zeus processor includes a register that effectively contains the current thread number that reads the register. In this way, threads running identical code can discover their own identity.

It is required that a sufficient number of bits be implemented so that each thread receives a distinct value. Values must be consecutive, unsigned and include a zero value. The remaining unimplemented bits must be zero whenever read. Writes to this register are ignored.



Physical address

The physical address of the Thread Register, byte b is:



Definition

```

def data ← AccessPhysicalThread(pa,op,wdata) as
  case op of
    R:
      data ← 064 11 Thread
    W:
      // nothing
  endcase
enddef

```

Index

A

Access detail required
 by global TB...96, 109, 119, 122, 125, 127, 130,
 132, 134, 287, 294, 302, 307, 310, 312, 315
 by local TB96, 109, 119, 122, 125, 127, 130, 132,
 134, 287, 293, 302, 307, 310, 312, 315
 by tag96, 108, 119, 122, 125, 127, 130, 132, 134,
 287, 293, 302, 307, 310, 312, 315
 Access disallowed
 by global TB...96, 108, 119, 122, 125, 127, 130,
 132, 134, 287, 293, 302, 307, 310, 312, 315
 by local TB96, 108, 119, 122, 125, 127, 130, 132,
 134, 287, 293, 302, 307, 310, 312, 315
 by tag96, 108, 119, 122, 125, 127, 130, 132, 134,
 287, 293, 302, 307, 310, 312, 315
 by virtual address94, 96, 108, 111, 116, 119, 122,
 125, 127, 130, 132, 134, 287, 293, 302, 307,
 310, 312, 315
 Address.....73
 add
 immediate.....80
 signed
 check overflow.....80
 unsigned
 check overflow.....80
 signed
 check overflow.....73
 unsigned
 check overflow.....73
 add.....73
 and.....73
 immediate.....80
 not.....73
 and not
 immediate.....80
 compare
 and
 equal zero.....76
 not equal zero.....76
 equal.....76
 zero.....76
 greater
 signed76
 unsigned.....76
 zero signed.....76
 greater equal
 signed76
 unsigned.....76
 zero signed.....76
 less
 signed76
 unsigned.....76

zero signed76
 less equal
 signed.....76
 unsigned.....76
 zero signed.....76
 not equal76
 zero76
 Compare.....76
 copy.....80
 copy immediate79
 Copy Immediate.....79
 exclusive nor.....73
 fixed point arithmetic exception76
 Immediate80
 Reversed.....83
 multiplex.....93
 negate.....83
 signed check overflow.....83
 no operation.....76
 not.....80
 not and.....73
 immediate ..80
 not or73
 immediate ..80
 or 73
 immediate ..80
 not.....73
 or not
 immediate ..80
 Reversed.....86
 set.....79
 and
 equal zero86
 not equal zero ..86
 equal86
 immediate ..83
 equal zero.....86
 greater
 immediate
 unsigned83
 signed86
 unsigned ..86
 zero signed ..86
 greater equal
 immediate
 signed83
 unsigned ..83
 signed86
 unsigned86
 zero signed ..86
 less
 immediate
 signed83
 unsigned ..83

signed	86
unsigned	86
zero signed	86
less equal	
immediate	
signed	83
unsigned	83
signed	86
unsigned	86
zero signed	86
not equal	
immediate	83
not equal zero	86
set and	
equal	
immediate	83
not equal	
immediate	83
Shift	
Left	
Immediate	
Add	89
Subtract	90
Shift Immediate	91
shift left	
immediate	91
add	89
signed check overflow	91
subtract	90
unsigned check overflow	91
shift right	
immediate	
unsigned	
signed	
shift right	
immediate	91
subtract	86
signed check overflow	86
unsigned check overflow	86
subtract immediate	83
signed	
check overflow	83
unsigned	
check overflow	83
Ternary	93
xnor	
immediate	80
xor	73
immediate	80
zero	79
Always reserved	72
Always Reserved	72
architecture description registers	407
Arithmetic Operations	39
Branch	21, 94
and	
equal zero	98
not equal zero	98
back	95
barrier	97
Conditional	98
Floating-Point	101
Visibility	103
down	105
equal	98
floating-point	
double	101
half	101
quad	101
single	101
equal zero	98
gateway	106
Gateway	106
greater	
floating-point	
double	101
half	101
quad	101
single	101
signed	98
unsigned	98
zero	
signed	98
greater equal	
floating-point	
double	101
half	101
quad	101
single	101
signed	98
unsigned	98
zero signed	98
halt	110
Hint	111
Immediate	112
immediate	113
link	114
Immediate	113
link	114
invisible	
floating-point	
single	103
less	
floating-point	
double	101
half	101
quad	101
single	101

unsigned	98	nibbles	206
zero		octlets	206
signed	98	pecks	206
less equal		quadlets	206
floating-point		unsigned	
double	101	bytes	206
half	101	doublets	206
quad	101	hexlet	206
single	101	nibbles	206
signed	98	octlets	206
unsigned	98	pecks	206
zero signed	98	quadlets	206
less greater		signed	
floating-point		bytes	187
double	101	doublets	187
half	101	hexlet	187
quad	101	nibbles	187
single	101	octlets	187
link	115	pecks	187
Link	115	quadlets	187
no operation	98	unsigned	
not equal	98	bytes	187
not equal zero	98	doublets	187
not invisible		hexlet	187
floating-point		nibbles	187
single	103	octlets	187
not visible		pecks	187
floating-point		quadlets	187
single	103	copy	207
signed		deposit	
less	98	merge	
visible		bits	201
floating-point		bytes	201
single	103	doublets	201
Branch back	95	hexlet	201
Branch barrier	97	nibbles	201
Branch Conditionally	22	octlets	201
Branch down	105	pecks	201
Branch halt	110	quadlets	201
C			
checkpoint	401, 402	signed	
Classical Pipeline Structures	49	bytes	196
Clock		doublets	196
Cycle	412	hexlet	196
Event	412	nibbles	196
Watchdog	412	octlets	196
Compare-set	23	pecks	196
control register	409, 410, 411	quadlets	196
Crossbar	187	expand	
compress		immediate	
immediate		signed	
signed		bytes	206
bytes	206	doublets	206
doublets	206	hexlet	206
hexlet	206	nibbles	206
		octlets	206
		pecks	206
		quadlets	206

unsigned	
bytes	206
doublets	206
hexlet	206
nibbles	206
octlets	206
pecks	206
quadlets	206
signed	
bytes	187
doublets	187
hexlet	187
nibbles	187
octlets	187
pecks	187
quadlets	187
unsigned	
bytes	187
doublets	187
hexlet	187
nibbles	187
octlets	187
pecks	187
quadlets	187
extend	
immediate	
signed	
bytes	197
doublets	197
hexlet	197
nibbles	197
octlets	197
pecks	197
quadlets	197
extract	192
Extract	192
Field	196
Field Inplace	201
Inplace	204
no operation	207
rotate	
left	
bytes	187
doublets	187
hexlet	187
nibbles	187
octlets	187
pecks	187
quadlets	187
right	
bytes	187
doublets	187
hexlet	188
nibbles	187
octlets	188
pecks	187
quadlets	187
rotate left	
immediate	
bytes	206
doublets	206
hexlet	206
nibbles	206
octlets	206
pecks	206
quadlets	206
rotate right	
immediate	
bytes	206
doublets	206
hexlet	207
nibbles	206
octlets	207
pecks	206
quadlets	207
select bytes	220
shift left	
bytes	188
doublets	188
hexlet	188
immediate	
bytes	207
doublets	207
hexlet	207
nibbles	207
octlets	207
pecks	207
quadlets	207
signed bytes	
check overflow	207
signed doublets	
check overflow	207
signed hexlet	
check overflow	207
signed nibbles	
check overflow	207
signed octlets	
check overflow	207
signed pecks	
check overflow	207
signed quadlets	
check overflow	207
unsigned bytes	
check overflow	207
unsigned doublets	
check overflow	207
unsigned hexlet	
check overflow	207
unsigned nibbles	
check overflow	207
unsigned octlets	
check overflow	207
unsigned pecks	
check overflow	207

unsigned quadlets	207
check overflow	207
merge	
bytes	204
doublets.....	204
hexlet.....	204
immediate	
bytes	211
doublets	211
hexlet	211
nibbles	211
octlets.....	211
pecks.....	211
quadlets.....	211
nibbles.....	204
octlets.....	204
pecks.....	204
quadlets.....	204
signed	
bytes	
check overflow	188
doublets	
check overflow	188
hexlet	
check overflow	188
nibbles	
check overflow	188
octlets	
check overflow	188
pecks	
check overflow	188
quadlets	
check overflow	188
unsigned	
bytes	
check overflow	188
doublets	
check overflow	188
hexlet	
check overflow	188
nibbles	
check overflow	188
octlets	
check overflow	188
pecks	
check overflow	188
quadlets	
check overflow	188
shift right	
immediate	
unsigned	
bytes.....	207
doublets.....	207
hexlet.....	207
nibbles	207
octlets	207
pecks	207
quadlets	207
bytes	204
doublets	204
hexlet.....	204
immediate	
bytes.....	211
doublets	211
hexlet	211
nibbles	211
octlets	211
pecks	211
quadlets	211
nibbles.....	204
octlets	204
pecks	204
quadlets	204
unsigned	
bytes	188
doublets	188
hexlet	188
nibbles	188
octlets	188
pecks	188
quadlets	188
Short Immediate.....	206
Inplace.....	211
shuffle	
within bytes	213
within doublets.....	213
within hexlet	213
within octlets	213
within pecks.....	213
within quadlets	213
within triclet.....	213
Shuffle.....	213
signed	
shift right	
bytes	188
doublets	188
hexlet.....	188
immediate	
bytes.....	207
doublets	207
hexlet	207
nibbles	207
octlets	207
pecks	207
quadlets	207
nibbles.....	188
octlets	188
pecks	188
quadlets	188

swizzle.....	219
Swizzle.....	219
Ternary.....	220
unsigned	
deposit	
bytes	196
doublets.....	196
hexlet.....	196
nibbles.....	196
octlets.....	196
pecks.....	196
quadlets.....	196
extend	
immediate	
bytes	197
doublets.....	197
hexlet	197
nibbles	197
octlets.....	197
pecks.....	197
quadlets.....	197
withdraw	
bytes	196
doublets.....	196
hexlet.....	196
nibbles.....	196
octlets.....	196
pecks.....	196
quadlets.....	196
withdraw	
bytes	196
doublets.....	196
hexlet.....	196
nibbles	196
octlets.....	196
pecks.....	196
quadlets.....	196
D	
Data-handling Operations	36
E	
Ensemble.....	221
absolute value	
floating-point	
double	276
half.....	276
quad.....	276
single.....	276
add	
floating-point	
double	258
half.....	258
quad.....	258
single.....	258
complex	
multiply	
bytes	221
doublets.....	221
quadlets.....	221
convert	
floating-point	
double from octlets.....	277
double from quad.....	276
double from single	277
doublets from half.....	278
half from doublets.....	276, 277
half from single.....	276
hexlet from quad	278
hexlet from quad	278
octlets from double.....	278
octlets from double.....	278
quad from double.....	277
quad from hexlet.....	277
quadlets from single	278
single from double	276
single from half.....	277
single from quadlets.....	277
floating-point	
single from double	276
convolve	
complex	
bytes	221
doublets	221
floating-point	
double	
big-endian.....	232
little-endian	232
half	
big-endian.....	232
little-endian	232
single	
big-endian.....	232
little-endian	232
quadlets.....	221
extract	
immediate	
mixed-signed bytes	
big-endian	
ceiling	226
floor.....	226
nearest	226
zero.....	226
little-endian	
ceiling	226
floor.....	226
nearest	226
zero	227
mixed-signed doublets	
big-endian	
ceiling	226
floor.....	226

nearest.....	226
zero	226
little-endian	
ceiling.....	227
floor	227
nearest.....	227
zero	227
mixed-signed octlets	
big-endian	
ceiling.....	226
floor	226
nearest.....	226
zero	226
little-endian	
ceiling.....	227
floor	227
nearest.....	227
zero	227
mixed-signed quadlets	
big-endian	
ceiling.....	226
floor	226
nearest.....	226
zero	226
little-endian	
ceiling.....	227
floor	227
nearest.....	227
zero	227
signed bytes	
big-endian	
ceiling.....	225
floor	225
nearest.....	225
zero	225
little-endian	
ceiling.....	226
floor	226
nearest.....	226
zero	226
signed complex bytes	
big-endian	
ceiling.....	225
floor	225
nearest.....	225
zero	225
little-endian	
ceiling.....	225
floor	225
nearest.....	225
zero	225
signed complex doublets	
big-endian	
ceiling.....	225
floor	225
nearest.....	225
zero	225
little-endian	
ceiling	225
floor	225
nearest	225
zero	225
signed complex octlets	
big-endian	
ceiling.....	225
floor	225
nearest	225
zero	225
little-endian	
ceiling	225
floor	225
nearest	225
zero	225
signed complex quadlets	
big-endian	
ceiling	225
floor	225
nearest	225
zero	225
little-endian	
ceiling	225
floor	225
nearest	225
zero	225
signed doublets	
big-endian	
ceiling	225
floor	226
nearest	226
zero	226
little-endian	
ceiling	226
floor	226
nearest	226
zero	226
signed octlets	
big-endian	
ceiling	226
floor	226
nearest	226
zero	226
little-endian	
ceiling	226
floor	226
nearest	226
zero	226
signed quadlets	
big-endian	
ceiling	226
floor	226
nearest	226
zero	226
little-endian	
ceiling	226

floor	226
nearest.....	226
zero	226
unsigned bytes	
big-endian	
ceiling.....	227
floor	227
nearest.....	227
little-endian	
ceiling.....	227
floor	227
nearest.....	227
unsigned doublets	
big-endian	
ceiling.....	227
floor	227
nearest.....	227
little-endian	
ceiling.....	227
floor	227
nearest.....	227
unsigned octlets	
big-endian	
ceiling.....	227
floor	227
nearest.....	227
little-endian	
ceiling.....	227
floor	227
nearest.....	227
unsigned quadlets	
big-endian	
ceiling.....	227
floor	227
nearest.....	227
little-endian	
ceiling.....	227
floor	227
nearest.....	227
floating-point double	
big-endian.....	232
little-endian.....	232
floating-point half	
big-endian.....	232
little-endian.....	232
floating-point single	
big-endian.....	232
little-endian.....	232
mixed-signed	
bytes	221
doublets.....	221
octlets	221
quadlets	221
signed	
bytes	221
doublets.....	221
octlets	221
quadlets.....	221
unsigned	
bytes	
ceiling.....	221
doublets.....	221
octlets.....	221
quadlets.....	221
Convolve	
Extract	
Immediate	225
Floating-Point	232
copy	
floating-point	
double	276
half	276
quad.....	276
single	276
divide	
floating-point	
double	258, 259
half	258
quad.....	259
single	258
signed	
octlets.....	221
unsigned	
octlets.....	221
Ensembleconvert	
floating-point	
doubletsfromhalf.....	278
octletsfromdouble.....	278
quadletsfromsingle	278
extract	236
immediate	
signed	
bytes	
ceiling.....	244
floor	244
nearest.....	244
zero	244
doublets	
ceiling.....	244
floor	244
nearest.....	244
zero	244
octlets	
ceiling.....	244
floor	244
nearest.....	244
zero	244
quadlets	
ceiling.....	244
floor	244
nearest.....	244
zero	244
unsigned	
bytes	
ceiling.....	244

floor.....	244
nearest.....	244
doublets	
ceiling.....	244
floor.....	244
nearest.....	244
octlets	
ceiling.....	244
floor.....	244
nearest.....	244
quadlets	
ceiling.....	244
floor.....	244
nearest.....	244
Extract.....	236
Immediate.....	244
Immediate Inplace.....	251
floating-point	
reciprocal square root estimate	
double	278
half.....	277
quad.....	278
single.....	277
Floating-point.....	258
Inplace.....	261
Floating-point.....	264
log of most significant bit	
signed	
bytes	274
doublets.....	274
hexlet.....	274
octlets.....	274
quadlets.....	274
unsigned	
bytes	274
doublets.....	274
hexlet.....	274
octlets.....	274
quadlets.....	274
multiply	
add	
extract	
immediate	
mixed-signed	
bytes	251
doublets.....	251
octlets.....	251
quadlets.....	251
signed	
bytes	251
complex bytes.....	251
complex doublets.....	251
complex octlets.....	251
complex quadlets.....	251
doublets	251, 252
octlets.....	252
quadlets.....	252
unsigned	
bytes	
ceiling	245
doublets.....	252
octlets	252
quadlets	252
complex	
floating-point	
double.....	259
half.....	259
single.....	259
extract.....	236
immediate	
mixed-signed	
bytes	
ceiling	245
floor.....	245
nearest	245
zero.....	245
doublets	
ceiling	245
floor.....	245
nearest	245
zero.....	245
octlets	
ceiling	245
floor.....	245
nearest	245
zero.....	245
quadlets	
ceiling	245
floor.....	245
nearest	245
zero.....	245
signed	
bytes	
ceiling	244
floor.....	244
nearest	244
zero.....	244
doublets	
ceiling	244
floor.....	244
nearest	244
zero.....	244
octlets	
ceiling	245
floor.....	245
nearest	245
zero.....	245
quadlets	
ceiling	244
floor.....	244
nearest	245
zero.....	245
unsigned	
bytes	
ceiling	245

floor	245
nearest.....	245
doublets	
ceiling.....	245
floor	245
nearest.....	245
octlets	
ceiling.....	246
floor	246
nearest.....	246
quadlets	
ceiling.....	245
floor	245
nearest.....	246
floating-point	
double	259
half.....	259
quad.....	259
single.....	259
mixed-signed	
bytes	221
doublets.....	221
octlets	221
quadlets	221
polynomial	
bytes	221
doublets.....	221
octlets	221
quadlets	221
signed	
bytes	221
doublets.....	221
octlets	221
quadlets	221
sum	
mixed-signed	
bytes	
doublets	
octlets	
quadlets	
signed	
bytes	
doublets	
octlets	
quadlets	
unsigned	
bytes	
doublets	
octlets	
quadlets	
unsigned	
bytes	222
doublets.....	222
octlets	222
quadlets	222
multiply add	
complex	
bytes and doublets	261
doublets and quadlets	261
quadlets and octlets.....	261
complex floating-point	
double	264
half	264
single	264
floating-point	
double	264
half	264
quad.....	264
single	264
mixed-signed	
bytes and doublets	261
doublets and quadlets	261
octlets and hexlet.....	261
quadlets and octlets.....	261
signed	
bytes and doublets	261
doublets and quadlets	261
octlets and hexlet.....	261
quadlets and octlets.....	261
unsigned	
bytes and doublets	261
doublets and quadlets	261
octlets and hexlet.....	261
quadlets and octlets.....	261
multiply extract	
immediate	
complex	
bytes	
ceiling.....	245
floor	245
nearest.....	245
zero	245
doublets	
ceiling.....	245
floor	245
nearest.....	245
zero	245
octlets	
ceiling.....	245
floor	245
nearest.....	245
zero	245
quadlets	
ceiling.....	245
floor	245
nearest.....	245
zero	245
multiply subtract	
complex	
bytes and doublets	261
doublets and quadlets	261
quadlets and octlets.....	261
complex floating-point	
double	264

half.....	264	half	279
single.....	264	quad.....	279
floating-point		single	279
double	264	signed	
half.....	264	bytes	274
quad.....	264	doublets	274
single.....	264	octlets	274
mixed-signed		quadlets	274
bytes and doublets.....	261	unsigned	
doublets and quadlets.....	261	bits	274
octlets and hexlet.....	261	bytes	274
quadlets and octlets	261	doublets	274
signed		octlets.....	274
bytes and doublets.....	261	quadlets.....	274
doublets and quadlets.....	261	Ternary	269
octlets and hexlet.....	261	Ternary	
quadlets and octlets	261	Floating-point	272
unsigned		Unary	274
bytes and doublets.....	261	Floating-point	276
doublets and quadlets.....	261	Ensemble absolute value	
octlets and hexlet.....	261	floating-point	
quadlets and octlets	261	half.....	276
multiply sum		Ensemble copy	
complex		floating-point	
bytes		half.....	276
doublets		Ensemble scale add	
quadlets		floating-point	
negate		double.....	272
floating-point		half.....	272
double	277	single.....	272
half.....	277	F	
quad.....	277	Fixed-point	21
single.....	277	Fixed-point arithmetic .75, 78, 82, 85, 88, 92, 137, 153, 162, 167, 175, 191, 210	
reciprocal estimate		Floating-point	22
floating-point		Floating-point arithmetic . 156, 178, 260, 268, 282, 307	
double	277	Forwarding	55
half.....	277	G	
quad.....	277	Galois Field Operations	39
single.....	277	Gateway	18
Reversed		Gateway disallowed	108
Floating-point		Global TB miss.. 96, 109, 119, 122, 125, 127, 134, 287, 294, 302, 307, 310, 312, 315	
scale add extract	236	Group	
square root		add	
floating-point		bits	141
double	279	bytes.....	135
half.....	278	doublets.....	135
quad.....	279	halve	
single.....	279	signed	
subtract		bytes	135
floating-point		ceiling	138
double	267		
half.....	267		
quad.....	267		
single.....	267		
sum			
floating-point			
double	279		

floor.....	138
nearest.....	138
zero.....	138
doublets	
ceiling.....	138
floor.....	138
nearest.....	138
zero.....	138
hexlet	
ceiling.....	138
floor.....	138
nearest.....	138
zero.....	138
octlets	
ceiling.....	138
floor.....	138
nearest.....	138
zero.....	138
quadlets	
ceiling.....	138
floor.....	138
nearest.....	138
zero.....	138
unsigned	
bytes	
ceiling.....	138
floor.....	138
nearest.....	138
doublets	
ceiling.....	138
floor.....	138
nearest.....	138
hexlet	
ceiling.....	138
floor.....	138
nearest.....	138
octlets	
ceiling.....	138
floor.....	138
nearest.....	138
quadlets	
ceiling.....	138
floor.....	138
nearest.....	138
hexlet.....	135
immediate	
doublet.....	159
hexlet.....	159
octlet.....	159
quadlet.....	159
signed	
doublet	
check overflow.....	159
hexlet	
check overflow.....	159
octlet	
check overflow.....	159
quadlets	
check overflow.....	159
quadlet	
check overflow.....	159
unsigned	
check overflow.....	159
doublet	
check overflow.....	159
hexlet	
check overflow.....	159
octlet	
check overflow.....	159
quadlet	
check overflow.....	159
limit	
signed	
bytes.....	135
doublets.....	135
hexlet.....	135
octlets.....	135
quadlets.....	135
unsigned	
bytes.....	135
doublets.....	135
hexlet.....	135
octlets.....	135
quadlets.....	135
octlets.....	135
quadlets.....	135
signed	
bytes	
check overflow.....	135
doublets	
check overflow.....	135
hexlet	
check overflow.....	135
octlets	
check overflow.....	135
quadlets	
check overflow.....	135
unsigned	
bytes	
check overflow.....	135
doublets	
check overflow.....	135
hexlet	
check overflow.....	135
octlets	
check overflow.....	135
quadlets	
check overflow.....	135
Add.....	135
Halve.....	138
add add add	
bits.....	141
bytes.....	168
doublets.....	168
hexlet.....	168
octlets.....	168
quadlets.....	168

add add subtract	
bits.....	141
bytes.....	168
doublets.....	168
hexlet.....	168
octlets.....	168
quadlets.....	168
add subtract add	
bytes.....	168
doublets.....	168
hexlet.....	168
octlets.....	168
quadlets.....	168
and	141
immediate	
doublet.....	159
hexlet.....	159
octlet.....	159
quadlet.....	159
and not	141
immediate	
doublet.....	160
hexlet.....	160
octlet.....	160
quadlet.....	160
boolean.....	141
Boolean.....	141
compare	
and	
equal zero	
bytes.....	148
doublets.....	148
hexlet	148
octlets.....	148
quadlets.....	148
not equal zero	
bytes.....	148
doublets.....	148
hexlet	148
octlets.....	148
quadlets.....	148
equal	
bytes	148
doublets.....	148
floating-point	
double.....	154
half	154
quad	154
single.....	154
hexlet.....	148
octlets	148
quadlets	148
zero	
signed	
bytes	150
doublets.....	150
hexlet.....	150
octlets.....	168
quadlets.....	168
greater	
floating-point	
double.....	155
half.....	155
quad.....	155
single.....	155
unsigned	
bytes.....	150
doublets.....	150
hexlet.....	150
octlets	150
quadlets	150
zero	
signed	
bytes	150
doublets	150
hexlet	150
octlets	150
quadlets	150
greater equal	
signed	
bytes.....	148
doublets.....	148
hexlet.....	148
octlets	148
quadlets	148
unsigned	
bytes.....	148
doublets.....	148
hexlet	148
octlets	148
quadlets	148
zero signed	
bytes.....	150
doublets.....	150
hexlet	150
octlets	150
quadlets	150
greater or equal	
floating-point	
double.....	154
half.....	154
quad.....	154
single.....	154
less	
floating-point	
double.....	154
half.....	154
quad.....	154
single.....	154
unsigned	
bytes.....	148
doublets.....	148
hexlet	148
octlets	148

quadlets.....	148
zero	
signed	
bytes	150
doublets.....	150
hexlet.....	150
octlets.....	150
quadlets.....	150
less equal	
floating-point	
double.....	155
half	155
quad	155
single.....	155
signed	
bytes.....	150
doublets.....	150
hexlet	150
octlets.....	150
quadlets.....	150
unsigned	
bytes.....	150
doublets.....	150
hexlet	150
octlets.....	150
quadlets.....	150
zero signed	
bytes	150
doublets.....	150
hexlet	151
octlets.....	150
quadlets.....	150
less or greater	
floating-point	
double.....	154
half	154
quad	154
single.....	154
not equal	
bytes	148
doublets.....	148
hexlet	149
octlets.....	148
quadlets.....	148
zero signed	
bytes	151
doublets.....	151
hexlet	151
octlets.....	151
quadlets.....	151
signed	
greater	
bytes	150
doublets.....	150
hexlet	150
octlets.....	150
quadlets.....	150
less	
bytes.....	148
doublets.....	148
hexlet.....	148
octlets	148
quadlets	148
Compare.....	148
Floating-point	
copy.....	141, 160
copy immediate	
byte	157
doublet	157
hexlet	157
octlet	157
quadlet	157
Copy Immediate.....	157
exclusive-nor.....	142
immediate	
doublet.....	160
hexlet	160
octlet	160
quadlet	160
exclusive-or.....	142
immediate	
doublet.....	159
hexlet	159
octlet	159
quadlet	159
fixed point arithmetic exception	151
Galois field multiply	
byte	269
octlet	269
Immediate	159
Reversed.....	163
Inplace	168
multiplex.....	186
nand	141
negate	
doublet	164
hexlet	164
octlet	164
quadlet	164
signed	
doublet	
check overflow.....	164
hexlet	
check overflow.....	164
octlet	
check overflow.....	164
quadlet	
check overflow.....	164
no operation.....	151
nor	141
not	141, 160
not and	
immediate	
doublet.....	159

hexlet.....	159
octlet.....	159
quadlet.....	159
not or	
immediate	
doublet	159
hexlet.....	159
octlet.....	159
quadlet.....	159
or 141	
immediate	
doublet	159
hexlet.....	159
octlet.....	159
quadlet.....	159
or not.....	141
immediate	
doublet	160
hexlet.....	160
octlet.....	160
quadlet.....	160
Reversed.....	170
Floating-point	
set.....	157
and	
equal zero	
bits.....	141
bytes.....	170
doublets.....	170
hexlet.....	170
octlets.....	170
quadlets.....	170
not equal zero	
bits.....	141
bytes.....	170
doublets.....	170
hexlet.....	170
octlets.....	170
quadlets.....	170
equal	
bits	141
bytes	170
doublets.....	170
floating-point	
double	176
half	176
quad	176
single	176
hexlet.....	170
octlets.....	170
quadlets.....	170
zero	
bytes	172
doublets	172
hexlet	172
octlets	172
quadlets	172
greater	
floating-point	
double.....	177
half.....	177
quad.....	177
single.....	177
immediate	
unsigned	
doublets	164
hexlet.....	164
octlets.....	164
quadlets.....	164
signed	
bits	141
unsigned	
bits	141
bytes.....	173
doublets.....	173
hexlet.....	173
octlets.....	173
quadlets.....	173
zero	
signed	
bits	141
bytes	172
doublets	172
hexlet	172
octlets	172
quadlets	172
greater equal	
floating-point	
double.....	176
half.....	176
quad	176
single	176
signed	
bits	141
bytes	170
doublets	170
hexlet	170
octlets	170
quadlets	170
unsigned	
bits	141
bytes	170
doublets	170
hexlet	170
octlets	170
quadlets	170
zero signed	
bits	141
bytes	172
doublets	172
hexlet	172
octlets	172
quadlets	172
immediate	

signed	
greater	
doublets.....	164
hexlet.....	164
octlets	164
quadlets.....	164
less	
floating-point	
double.....	176
half.....	176
quad.....	176
single.....	176
signed	
bits.....	141
unsigned	
bits.....	141
bytes.....	170
doublets.....	170
hexlet.....	170
octlets.....	170
quadlets.....	170
zero	
signed	
bits.....	141
bytes	172
doublets.....	172
hexlet.....	172
octlets.....	172
quadlets.....	172
less equal	
floating-point	
double.....	177
half	177
quad.....	177
single.....	177
immediate	
signed	
doublets.....	164
hexlet.....	164
octlets	164
quadlets.....	164
unsigned	
doublets.....	164
hexlet.....	164
octlets	164
quadlets.....	164
signed	
bits.....	141
bytes	172
doublets.....	172
hexlet.....	172
octlets.....	172
quadlets.....	172
signed zero	
bytes	172
doublets	172
octlets.....	172
quadlets	172
unsigned	
bits	141
bytes	172
doublets.....	172
hexlet.....	172
octlets	172
quadlets	172
zero signed	
bits	141
hexlet.....	172
less greater	
floating-point	
double.....	176
half	176
quad.....	176
single.....	176
not equal	
bits.....	141
bytes	170
doublets.....	170
hexlet.....	171
octlets	170
quadlets	170
zero	
bytes	172
doublets	172
hexlet	172
octlets	172
quadlets	172
signed	
greater	
bytes	172
doublets	172
hexlet	173
octlets	172
quadlets	172
less	
bytes	170
doublets	170
hexlet	170
octlets	170
quadlets	170
set	141
set and	
equal zero	
immediate	
doublets.....	163
hexlet	163
octlets	163
quadlets	163
not equal zero	
immediate	
doublets	163
hexlet	163
octlets	163
quadlets	163

big-endian:	120
little-endian:	120
big-endian:	120
little-endian:	120
quadlet	
aligned	
big-endian:	120
little-endian:	120
big-endian:	120
little-endian:	120
unsigned	
byte:	120
doublet	
aligned	
big-endian:	120
little-endian:	120
big-endian:	120
little-endian:	120
octet	
aligned	
big-endian:	120
little-endian:	120
big-endian:	120
little-endian:	120
quadlet	
aligned	
big-endian:	120
little-endian:	120
big-endian:	120
little-endian:	120
Immediate.....	120
signed	
byte.....	117
doublet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
octet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
quadlet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
unsigned	
byte.....	117
doublet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
octet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
big-endian	117
little-endian	117
quadlet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
octlet	
aligned	
big-endian.....	117
little-endian	117
big-endian	117
little-endian	117
Load and Store.....	21
Local TB miss	96, 109, 119, 122, 125, 127, 130, 132, 134, 287, 294, 302, 307, 310, 312, 315
M	
machine check.....	400
Memory Management.....	316
most-privileged level.....	316
multiprocessor.....	316
P	
Pipeline Organization.....	49
Prediction.....	54
Procedure Calling Conventions.....	40
R	
Register Usage.....	40
Reserved instruction.....	200, 203
Reserved Instruction	94, 96, 97, 105, 108, 110, 116, 158, 210, 212, 218
reset.....	399
Resources.....	55
Rounding	24
S	
set on compare.....	35, 54
Software Conventions.....	40
start vector address.....	399, 400, 402
status register.....	407, 408
Store	123
add swap immediate	
octlet	
aligned big-endian	131
aligned little-endian	131
add swap octlet	
aligned big-endian	133
aligned little-endian	133
byte.....	123
compare swap immediate	

octlet	
aligned big-endian.....	131
aligned little-endian	
immediate.....	131
compare swap octlet	
aligned big-endian	133
aligned little-endian.....	133
double	
aligned	
big-endian	123
little-endian.....	123
big-endian.....	123
little-endian	123
double compare swap octlet	
aligned big-endian	126
aligned little-endian.....	126
hexlet	
aligned	
big-endian.....	123
little-endian.....	123
big-endian.....	123
little-endian	123
immediate	
byte.....	128
double	
aligned	
big-endian.....	128
little-endian	128
big-endian	128
little-endian.....	128
hexlet	
aligned	
big-endian.....	128
little-endian	128
big-endian	128
little-endian	128
octlet	
aligned	
big-endian.....	128
little-endian	128
big-endian	128
little-endian	128
quadlet	
aligned	
big-endian.....	128
little-endian	128
big-endian	128
little-endian	128
Immediate.....	128
Inplace	131
Inplace.....	133
multiplex immediate	
octlet	
aligned big-endian.....	128
aligned little-endian	128
multiplex octlet	
aligned big-endian	123
aligned little-endian.....	123
multiplex swap immediate	
octlet	
aligned big-endian	131
aligned little-endian	131
multiplex swap octlet	
aligned big-endian.....	133
aligned little-endian.....	133
octlet	
aligned	
big-endian.....	123
little-endian	123
big-endian	123
little-endian	123
quadlet	
aligned	
big-endian.....	123
little-endian	123
big-endian	123
little-endian	123
Store Double Compare Swap	126
Superspring Pipeline.....	51
Superstring Pipeline.....	50
Superthread Pipeline	52

W

Wide	
multiply	
matrix	
extract	
big-endian:	288
immediate	
complex	
bytes.....	295
doublets.....	296
octlets	296
quadlets	296
signed	
bytes.....	295
doublets.....	295
octlets	295
quadlets	295
unsigned	
bytes.....	297
doublets.....	297
octlets	297
quadlets	297
little-endian:	288
Galois big-endian	308
Galois little-endian	308
mixed-signed	
byte	
big-endian.....	283
little-endian	283
doublet	
big-endian.....	283

little-endian.....	283
quadlet	
big-endian.....	283
little-endian.....	283
polynomial byte	
big-endian.....	283
little-endian.....	283
polynomial doublet	
big-endian.....	283
little-endian.....	283
polynomial quadlet	
big-endian.....	283
little-endian.....	283
signed	
byte	
big-endian.....	283
little-endian.....	283
complex byte	
big-endian.....	283
little-endian.....	283
complex doublet	
big-endian.....	283
little-endian.....	283
doublet	
big-endian.....	283
little-endian.....	283
quadlet	
big-endian.....	283
little-endian.....	283
unsigned	
byte	
big-endian.....	283
little-endian.....	283
doublet	
big-endian.....	283
little-endian.....	283
quadlet	
big-endian.....	283
little-endian.....	283
Multiply	
Matrix	283
Extract.....	288
Immediate	295
Floating-Point	303
Galois	308
multiply matrix	
complex	
floating-point	
double	
big-endian.....	303
little-endian.....	303
half	
big-endian.....	303
little-endian.....	303
single	
big-endian.....	303
little-endian.....	303
extract immediate mixed-signed bytes	
big-endian	
ceiling.....	296
floor	296
nearest.....	296
zero	296
little-endian	
ceiling.....	296
floor	296
nearest.....	296
zero	296
doublets	
big-endian	
ceiling.....	296
floor	296
nearest.....	296
zero	296
little-endian	
ceiling.....	296
floor	296
nearest.....	296
zero	296
octlets	
big-endian	
ceiling.....	296
floor	297
nearest.....	297
zero	297
little-endian	
ceiling.....	296
floor	297
nearest.....	297
zero	297
quadlets	
big-endian	
ceiling.....	296
floor	296
nearest.....	296
zero	296
little-endian	
ceiling.....	296
floor	296
nearest.....	296
zero	296
floating-point	
double	
big-endian	303
little-endian	303
half	
big-endian	303
little-endian	303
single	
big-endian	303
little-endian	303
switch	
big-endian	311

little-endian	311	little-endian	313
Switch.....	311	octlets	
translate		big-endian.....	313
bytes		little-endian	313
big-endian	313	quadlets	
little-endian.....	313	big-endian.....	313
doublets		little-endian	313
big-endian	313	Translate	313

REGISTERED
MICROUNITY
DO NOT REPRODUCE
CONFIDENTIAL
COPY #0
Copyright